

Université Libre de Bruxelles

*Institut de Recherches Interdisciplinaires
et de Développements en Intelligence Artificielle*

**Property-driven Design for Robot Swarms:
A Design Method Based on Prescriptive
Modeling and Model Checking**

M. BRAMBILLA, M. DORIGO, and M. BIRATTARI

IRIDIA – Technical Report Series

Technical Report No.
TR/IRIDIA/2014-003

January 2014

IRIDIA – Technical Report Series
ISSN 1781-3794

Published by:

IRIDIA, *Institut de Recherches Interdisciplinaires
et de Développements en Intelligence Artificielle*
UNIVERSITÉ LIBRE DE BRUXELLES
Av F. D. Roosevelt 50, CP 194/6
1050 Bruxelles, Belgium

Technical report number TR/IRIDIA/2014-003

The information provided is the sole responsibility of the authors and does not necessarily reflect the opinion of the members of IRIDIA. The authors take full responsibility for any copyright breaches that may result from publication of this paper in the IRIDIA – Technical Report Series. IRIDIA is not responsible for any use that might be made of data appearing in this publication.

Property-driven design for robot swarms: A design method based on prescriptive modeling and model checking

MANUELE BRAMBILLA, MARCO DORIGO, and MAURO BIRATTARI, IRIDIA, Université Libre de Bruxelles, Brussels, Belgium

In this paper, we present property-driven design, a novel top-down design method for robot swarms based on prescriptive modeling and model checking. Traditionally robot swarms have been developed using a code-and-fix approach: in a bottom-up iterative process, the developer tests and improves the individual behaviors of the robots until the desired collective behavior is obtained. The code-and-fix approach is unstructured and the quality of the obtained swarm depends completely on the expertise and ingenuity of the developer. Property-driven design is composed of four phases: first, the developer formally specifies the requirements of the robot swarm by stating its desired properties; second, the developer creates a prescriptive model of the swarm and uses model checking to verify that this prescriptive model satisfies the desired properties; third, using the prescriptive model as a blueprint, the developer implements a simulated version of the desired robot swarm and validates the prescriptive model developed in the previous steps; fourth, the developer implements the desired robot swarm and validates the previous steps. We demonstrate property-driven design using two case studies: aggregation and foraging.

Categories and Subject Descriptors: I.2.9 [Artificial Intelligence]: Robotics

General Terms: Design, Reliability, Verification

Additional Key Words and Phrases: Swarm robotics; Top-down design; Prescriptive modeling; Model checking; Aggregation; Foraging

ACM Reference Format:

Manuele Brambilla, Marco Dorigo, and Mauro Birattari, 2014. Property-driven design for robot swarms: A design method based on prescriptive modeling and model checking. *ACM Trans. Embedd. Comput. Syst.* V, N, Article A (January YYYY), 26 pages.

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

1. INTRODUCTION

Swarm robotics is an approach to the coordination of large groups of robots that takes inspiration from social insects, such as ants, bees and termites [Şahin 2005]. Swarm robotics aims at developing systems that are *fault tolerant*, *scalable* and *flexible* [Dorigo et al. 2014].

The research leading to the results presented in this paper has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement n.246939 and by the EU project ASCENS, 257414.

Manuele Brambilla, Marco Dorigo and Mauro Birattari acknowledge support from the F.R.S.-FNRS of Belgium's Wallonia-Brussels Federation, of which they are a F.R.I.A. Research Fellow, a Research Director and a Research Associate, respectively.

Author's addresses: M. Brambilla, M. Dorigo, M. Birattari, IRIDIA, CoDE, Université Libre de Bruxelles, 50 Av. Franklin Roosevelt CP 194/6, 1050 Brussels, Belgium; email: mbrambil@ulb.ac.be; mdorigo@ulb.ac.be; mbiro@ulb.ac.be.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© YYYY ACM 1539-9087/YYYY/01-ARTA \$15.00

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

Robot swarms are self-organized systems that can be observed at two levels: the individual, also called microscopic, level; and the collective, also called macroscopic, level. The individual level is the behavior displayed by a single robot. The collective level is the behavior displayed by the swarm and it is the result of the interaction of the individual behaviors.

On the one hand, this dual nature of swarm robotics systems is key in achieving fault tolerance, scalability and flexibility. On the other hand, it is the source of difficult design challenges. In fact, the swarm robotics engineer must *think at the collective-level, but develop at the individual-level*: developers of robot swarms are caught between collective-level missions, such as “monitor the perimeter of a building for intruders” or “carry these heavy objects from here to there”, and individual-level software, as the only controllable/programmable components of a robot swarm are the individual behaviors of the robots. Conversely, at the individual level, collective-level goals could be meaningless: for example, for a single robot it is impossible to monitor an entire building at the same time or transport an object if it is too heavy to be moved. Thus the developer needs to design the behavior of the individual robots so that their interaction will result in the collective-level behavior that is needed to accomplish the mission.

Unfortunately, the design and development of individual-level behaviors to obtain a desired swarm-level goal is, in general, very difficult, as it is difficult to predict and thus design the non-linear interactions of tens or hundreds individual robots that result in a desired collective behavior. The difficulty to predict and design such interactions and the lack of a centralized controller make traditional system engineering approaches ineffective [Wooldridge and Jennings 1998; Banzhaf and Pillay 2007].

Some approaches to the design of robot swarms have been proposed in the last years. However, as discussed in Section 2, these approaches present limitations and an effective approach to the top-down design of robot swarms is still missing.

In this paper, we present property-driven design, a novel top-down design method for robot swarms based on prescriptive modeling and model checking. The developer creates a prescriptive model of the desired robot swarm and uses it as a blueprint for the implementation and improvement of the final swarm. The use of model checking allows the developer to formally verify properties directly on the model, reducing the need for testing in simulation or with robots. In property-driven design, different “views” of the system to realize are produced, from the most abstract (the properties of the system) to the most concrete (the final robot swarm). This is similar to model-driven engineering [Miller and Mukerji 2003] where software is designed through a series of model transformations from platform-independent models to executable platform-specific models.

Property-driven design addresses the shortcomings of the existing approaches.

- It aims at providing a method to formally specify the requirements of the desired robot swarm;
- It reduces the risk of developing the “wrong” robot swarm, that is, a robot swarm that does not satisfy the requirements;
- It promotes the re-use of available models and tested solutions;
- It can be used to develop platform-independent models that help in identifying the best robotic platform to use;
- It helps to shift the focus of the development process from implementation to design.

Property-driven design is a step forward in the development of *swarm engineering*: the systematic application of scientific and technical knowledge to specify requirements, design, realize, verify, validate, operate and maintain an artificial swarm intelligence system [Brambilla et al. 2013].

To illustrate and validate property-driven design, we apply it to two case studies: aggregation and foraging.

In Section 2, we present the related literature on design methods and model checking for swarm robotics. In Section 3, we present property-driven design. In Section 4, we present the two case studies.

2. RELATED WORK

In this section we first discuss the literature on design methods and then the literature on model checking in swarm robotics.

Design methods. The design of multi-robot systems has been addressed in many research papers [Zambonelli et al. 2001; Bordini 2009; Goldberg and Mataric 2001]. However, the design of robot swarms poses challenges that are not present in other multi-robot systems. Indeed, the characteristics of robot swarms, such as high number of individuals, strong decentralization, simple behaviors, local communication and action, are usually regarded as characteristics that make a multi-robot system “too complex to manage effectively” [Wooldridge and Jennings 1998].

Traditional multi-robot approaches are thus of limited use when developing robot swarms. For this reason other ad-hoc design approaches have been proposed.

Kazadi et al. [2009] developed a design approach based on Hamiltonian vector fields called the *Hamiltonian method*: starting from a mathematical description of a collective behavior, the method can be used to derive microscopic rules that minimize or maximize a selected numerical value (e.g., the virtual potential energy of a particular state of the swarm). The Hamiltonian method has the major drawback that it deals only with spatially-organizing behaviors such as pattern formation.

Berman et al. [2009] proposed a top-down approach to the design of a task allocation behavior. The authors describe the system as a Markov chain in which states represent tasks and edges represent the possibility for a robot to move from a task to another. Using a stochastic optimization method, it is possible to derive the probabilities that govern how robots change task in order to minimize the time needed to converge to the desired allocation. This approach is specific for task allocation and it has not been extended to other collective behaviors.

Hamann and Wörn [2008] proposed a method inspired by statistical physics. The authors use Langevin equations to describe the individual behaviors of the robots and, through analytical means, they derive a Fokker-Planck equation describing the collective behavior of the system. A similar approach was adopted also by Berman et al. [2011], who used a set of advection-diffusion-reaction partial differential equations to derive the individual behaviors of a swarm performing task allocation. Both methods are based on advanced mathematical techniques and on the ability for the developer to model the robot interactions. Moreover, such methods rely on ordinary or partial differential equations, which provide reliable results only if it is assumed that the swarm size tends to infinite. In swarm robotics this is very often not the case, since typically robot swarms are composed of no more than a hundred robots and often of just a few tens of robots [Brambilla et al. 2013].

Model checking in swarm robotics. Property-driven design is based on model checking [McMillan 1993], a technique to prove properties of a system in a formal way.

Dixon et al. [2012] were the first to apply model checking to swarm robotics. The authors used a microscopic model, in the form of the *and*-composition of individual-level models, and linear temporal logic to define properties of individual robots and of the swarm. This approach is not scalable, as the number of states of the model increases exponentially with the number of robots. Furthermore, linear temporal logic,

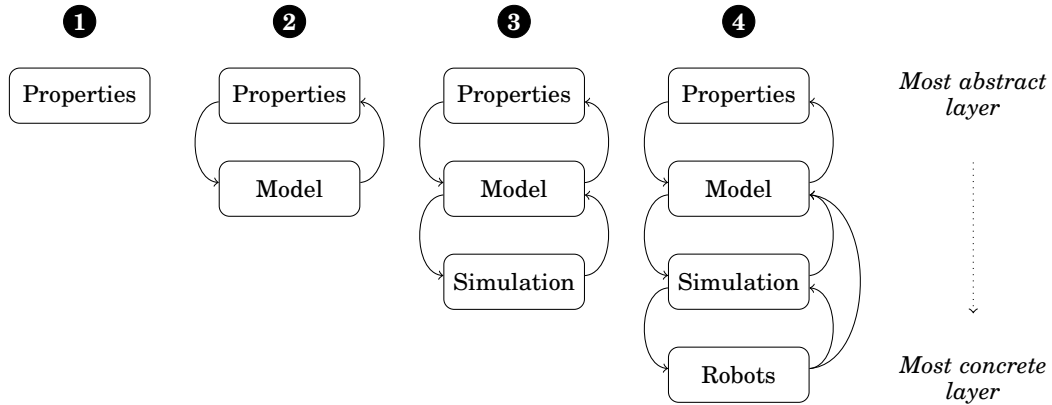


Fig. 1. The four phases of property-driven design.

which deals only with binary values (true/false) is a suboptimal choice to analyze robot swarms, which are systems characterized by stochastic properties.

Recently, Konur et al. [2012] adopted a different approach. The authors used model checking on a macroscopic model of a robot swarm performing foraging. They specified the desired properties of the system using *probabilistic computation tree logic* (PCTL), a temporal logic that includes probabilistic aspects (see Appendix I for a description of PCTL). This approach is able to overcome the limits of linear temporal logical. Moreover the use of a macroscopic model, instead of a microscopic one, allows this approach to deal with systems composed of tens of robots. We used a similar approach in a previous work [Brambilla et al. 2012] in which model checking and PCTL were used to verify properties of a robot swarm performing aggregation.

In a work on the use of Bio-PEPA in swarm robotics [Massink et al. 2013], we were the first to use statistical model checking (presented in Appendix I) to analyze a collective decision-making behavior. Statistical model checking overcomes the scalability issues of complete model checking allowing us to analyze models of large swarms.

3. PROPERTY DRIVEN DESIGN

Property-driven design is composed of four-phases: i) the requirements of the robot swarm are first formally described in the form of desired properties; ii) subsequently, a prescriptive model of the robot swarm is created; iii) this prescriptive model is used as a blueprint to implement and improve a simulated version of the desired robot swarm; iv) the final robot swarm is implemented.

A schema showing the different phases of property-driven design is presented in Figure 1.

In each of the phases of property-driven design, a new layer is added to the system. Layers differ in their level of abstraction: the *properties* layer is the most abstract, in which only the goal characteristics of the robot swarm are stated; the *robots* layer is the most concrete, in which the actual software for the real robots is developed and deployed. The addition of a new layer brings the system closer to its final state.

Each phase of property-driven design is characterized by a *development/validation* cycle: the focus of the developer is on the newly introduced layer, but all previously developed layers are still active, that is, they are still improved and expanded, should this be needed in order to guarantee the consistency of all layers. The newly introduced

layer provides the developer with further information on the system. This information is used to improve the system being developed, to validate its prescriptive model, and to verify its properties. For example, the development of the system in simulation provides the developer with new data that can be used to improve and validate the prescriptive model and further verify that the desired properties hold.

Phase one: Properties – In this phase, the developer formally specifies the requirements of the robot swarm in the form of desired properties. These properties are the distinguishing features of the robot swarm that the developer wants to realize. They can be task specific, such as *the system eventually completes task X*, or they can express more generic properties, such as *the system keeps working as long as there are at least N robots* or *the system will never be in state Y for more than t time-steps*. The clearer and more complete these properties are in this phase, the more the developed robot swarm will meet expectations. Clearly stated requirements help reducing the risk of developing “the wrong robot swarm.” For simplicity, we assume that requirements do not change during the development of the robot swarm.

Phase two: Model – In this phase, the developer creates a prescriptive model of the robot swarm. Usually, the prescriptive model describes how robots change state over time, where a state is an abstract simplified description of the actions of a robot (see also Appendix I). The prescriptive model should be sufficiently detailed to capture the behavior of the robots and their interaction, but should not be too detailed, in order to avoid unnecessary complication.

Once a first draft of the prescriptive model is produced, the desired properties stated in phase one are verified using model checking. As in test-driven development [Beck 2003], at first it is possible that the prescriptive model does not satisfy all the desired properties. In an iterative process, the developer expands and improves the prescriptive model, until the properties are satisfied. The outcome of this process is a prescriptive model of the collective behavior of the robot swarm that satisfies the stated properties.

Phase three: Simulation – In this phase, the developer uses the prescriptive model as a blueprint to implement and improve the robot swarm using a physics-based computer simulation (henceforth simply simulation). By blueprint we mean that the prescriptive model is used to identify the most relevant aspects of the robot swarm to realize. This allows the developer to focus on these aspects and neglect other minor details. For example, if a prescriptive model shows that, by entering state i , an individual robot affects the performance of the whole swarm more than by entering state j , the developer can focus on the first and temporarily ignore the second. Moreover, concentrating on the prescriptive model at design time allows the developer to direct his efforts towards high-level decisions rather than on the implementation.

It is possible that the implementation choices or other unforeseen aspects of the system yield in a simulated system that does not behave as predicted by the prescriptive model. In this case the developer must go back to the previous phases, modify the prescriptive model to consider the results obtained from the simulation, and verify whether the required properties still hold true.

Phase four: Robots – In the last phase, the developer realizes the final robot swarm. Similarly to the transition between the prescriptive model and the simulation, if the implementation on robots reveals that some assumptions made during the previous phases do not hold, it might be necessary to modify the simulated version or the prescriptive model, in order to keep all levels consistent.

4. CASE STUDIES

In this section, we illustrate property-driven design using two very common case studies from the swarm robotics literature [Brambilla et al. 2013]: aggregation and foraging.

In both case studies, we perform model checking using PRISM, a state-of-the-art suite for model checking [Kwiatkowska et al. 2004]. PRISM is free and is released as open source software under the GNU General Public License (GPL).¹

4.1. Aggregation

In the first case study, we tackle *aggregation*: robots have to cluster in an area of the environment. The robots have neither knowledge of the position of the other robots nor a map of the environment. We choose aggregation as a case study for various reasons: i) aggregation is a simple case study and this allows us to focus on the development process; ii) aggregation is a common case study in swarm robotics [Brambilla et al. 2013]; iii) aggregation possesses many of the salient traits of swarm robotics; it is completely distributed, it is based on simple robot-to-robot interactions, and it is characterized by stochasticity and spatial aspects.

The aggregation case study that we discuss in this paper is similar to the one presented by Jeanson et al. [2005]. We consider a dodecagonal environment with two black spots of equal size called *area A* and *area B*. We call *area C* the remaining white area. Each of the black spots is large enough to host all the robots. See Figure 5 for a picture of the environment. We consider three swarm sizes: 10, 20 and 50. We use three different arenas for the three different group sizes, respectively of 4.91 m², 19.63 m² and 50.26 m².

In the following, we will apply the 4-phase process explained in Section 3.

Phase one: Properties – The main property that the robot swarm must satisfy is “*eventually all the robots form an aggregate*”. The robots should aggregate as fast as possible either on area A or area B. We set a time limit of 1000 seconds. Using PRISM syntax, we can define the following property:

$$P \geq k \ [F \leq 1000 \ (a = N_t) \mid (b = N_t)] \quad (1)$$

In less formal terms, we want to know whether, in the first thousand seconds ($F \leq 1000$), the number of robots in area A or in area B is equal to the total number of robots in the swarm ($(a = N_t) \mid (b = N_t)$), with a probability greater or equal to k ($P \geq k$). The value of k depends on the size of the swarm: $k = 0.80$ for $N_t = 10$; $k = 0.40$ for $N_t = 20$; and $k = 0.01$ for $N_t = 50$;

Another property is that the aggregate, once formed, is stable for at least 10 seconds, that is, robots do not change state once the aggregate is formed. We want this to happen more than two thirds of the time an aggregate is formed:

$$(a = N_t) \mid (b = N_t) \Rightarrow P \geq 0.67 \ [G \geq 10 \ (a = N_t) \mid (b = N_t)] \quad (2)$$

In natural language, Property 2 can be expressed in this way: from the aggregate state $((a = N_t) \mid (b = N_t))$ is it true with probability of at least 0.67 ($\Rightarrow P \geq 0.67$) that the robot swarm stays for at least 10 seconds ($G \geq 10$) in the aggregate state?

Phase two: Model – To develop the prescriptive model for the aggregation case study, we consider the three areas in which the environment is divided. We define three states: S_a , S_b and S_c . A robot in area A or B is in state S_a or S_b , respectively. Robots

¹<http://www.prismmodelchecker.org>




Fig. 2. The prescriptive model for aggregation. Each state is used to count the number of robots in the corresponding area.

outside area A or B are in state S_c . We develop a discrete-time macroscopic prescriptive model. In a macroscopic model each state is associated with a counter to track the number of robots currently in that state (see Appendix I). In this model, we have three counters, a , b and c , where $a+b+c=N_t$, associated to the respective states. See Figure 2 for the model of the system.

In this initial stage of the definition of the prescriptive model, we assume that the system can be effectively described by a *non-spatial* model, that is, a model in which the trajectories of the robots are ignored and a robot can move instantaneously from area C to area A or B, and vice versa. Moreover, for the moment, we also ignore the effects of interferences between robots [Lerman et al. 2005]. In case these assumptions prove to be not realistic and the results obtained with the prescriptive model do not match those obtained in simulation or with the final robot swarm, we will modify them in the following phases, as explained in Section 3.

The first design attempt is the following: a robot performs random walk and, when it finds a black area, it stops. A robot stopped on a black area has a fixed probability to leave.

Since the prescriptive model is non-spatial and ignores interference, we consider only the geometric properties of the areas to compute p_{ca} , that is, the instantaneous probability that a robot transitions from S_c to S_a . A robot in area C can either go to area A, go to area B or stay in area C. This means that a robot in area C has a probability of going from area C to area A equal to $p_{ca} = \frac{A_A}{A_{arena}}$, of going from area C to area B equal to $p_{cb} = \frac{A_B}{A_{arena}}$, and of staying in area C equal to $p_{cc} = \frac{A_C}{A_{arena}} = 1 - (p_{ca} + p_{cb})$. Note that $p_{ca} = p_{cb}$, since the two areas have the same size.

The remaining probabilities depend on the behavior of the robots. The aggregate can be obtained in area A or area B, thus we set the probabilities of leaving these two areas to be equal: $p_{ac} = p_{bc}$. A robot in area A can only go to area C or stay in area A, thus $p_{aa} = 1 - p_{ac}$. The same holds for area B. From the above, it follows that $p_{aa} = p_{bb}$. The only independent probability remaining is p_{ac} . Through model checking, we can find the value of p_{ac} that maximizes the probability involved in the definition of Property 1.

Using model checking, we can find the best values for parameter p_{ac} and whether the required properties are satisfied. Using PRISM we can also compute the exact probabilities involved in the definition of the properties. Table I shows that this first attempt at tackling the aggregation case study is unsuccessful. The behavior obtains poor results and the system does not cope well with increasing group sizes.

An analysis of the prescriptive model can help us in improving the developed behavior. From the obtained results we observed that a fixed p_{ac} does not promote the formation of a single aggregate. A better solution is to let a robot decide whether to leave according

Table I. Model checking results for the first solution with a fixed p_{ac} . Column p_{ac} shows the best value of p_{ac} . Columns Pr 1 and Pr 2 show whether Property 1 and 2 are satisfied and the exact values of the probabilities involved in their definition.

N_t	A_A	A_{arena}	p_{ca}	p_{ac}	Pr 1	Pr 2
10	0.38 m ²	4.91 m ²	0.08	0.05	X (0.75)	X (0.46)
20	0.78 m ²	19.63 m ²	0.06	0.04	X (0.15)	X (0.07)
50	3.14 m ²	50.26 m ²	0.06	0.04	X (8.8×10^{-5})	X (3.7×10^{-5})

Table II. Model checking results for the second solution where $p_{ac} = 1 - p_{min-ac} * (N_s + 1)$. Column p_{min-ac} shows the best value of p_{min-ac} . Column Pr 1 and Pr 2 are defined as in Table I.

N_t	A_A	A_{arena}	p_{ca}	p_{min-ac}	Pr 1	Pr 2
10	0.38 m ²	4.91 m ²	0.08	[0.19, 0.24]	✓ (0.95)	✓ (0.92)
20	0.78 m ²	19.63 m ²	0.06	0.12	✓ (0.79)	✓ (0.87)
50	3.14 m ²	50.26 m ²	0.06	0.10	✓ (0.25)	✓ (0.71)

to the number of sensed robots around it: with only few robots nearby, the probability to leave the aggregate p_{ac} is high and vice versa. We set $p_{ac} = 1 - p_{min-ac} * (N_s + 1)$, where p_{min-ac} is the minimum staying probability we want for a robot and N_s is the number of other robots sensed. We add 1 to the number of robots sensed, as we include also the robot that is choosing its next action. Subsequently, using model checking, we find the best value of p_{min-ac} for the different group sizes. As reported in Table II, results are significantly better both for Property 1 and Property 2.

With the current prescriptive model we are also able to define specifications of the hardware capabilities of the robots: a ground sensor, to differentiate between the two black areas A and B and the white area C; a sensor to detect nearby robots; and wheels to move. An example of such a robot is the e-puck [Mondada et al. 2009], which can be extended with a range and bearing board that allows it to perceive the presence of neighboring robots [Gutiérrez et al. 2009].

Phase three: Simulation – In this aggregation case study, the prescriptive model captures well the microscopic behavior of the single robots, thus it is quite straightforward to implement the robot swarm in simulation. However, several implementation details are not explicitly present in the prescriptive model, such as how the robots perform random walk, and have now to be programmed explicitly.

We implement the robot swarm using the ARGoS simulator [Pinciroli et al. 2012]. Figure 3 presents a screenshot of the simulated robot swarm.

We perform three different sets of experiments, one for each group size. To validate the prescriptive model we measure the average time necessary to form a complete aggregate on 100 runs with different values of p_{min-ac} . The robots are deployed in a random position at the beginning of each experiment. Each experiment is halted when a complete aggregate is formed or after 10,000 seconds.

As reported in Figure 4, for all the three group sizes, the best results are obtained with the value p_{min-ac} predicted using the prescriptive model. However, the results related to Property 1 obtained with the simulated version of the robot swarm are usually worse than those predicted by the prescriptive model, in particular with 20 and 50 robots. With 10 robots and $p_{min-ac} = 0.22$ the simulated robot swarm was able to form a complete aggregate before 10,000 seconds 100 times out of 100, in line with the predictions of the prescriptive model. However, with 20 robots and $p_{min-ac} = 0.12$, an aggregate was formed in less than 1,000 seconds only 53 times out of 100, whereas in the prescriptive model this happened with probability 0.79. With 50 robots and




Fig. 3. A screenshot of the simulated version of the robot swarm with 20 robots.

Table III. A comparison between model checking and simulation. The table presents the probability involved in the definition of Property 1 (model checking) compared to the experimental results over 100 runs (simulation).

N_t	Model checking (old p_{ca} values)	Model checking (new p_{ca} values)	Simulation
10	0.95 with $p_{ca} = 0.08$	0.95 with $p_{ca} = 0.08$	100/100
20	0.79 with $p_{ca} = 0.06$	0.53 with $p_{ca} = 0.05$	53/100
50	0.25 with $p_{ca} = 0.06$	0.01 with $p_{ca} = 0.04$	2/100

$p_{min-ac} = 0.10$ the difference is even more evident: only 2 runs out of 100 resulted in an aggregation time of under 1,000 seconds whereas the prescriptive model predicted a probability of 0.25.

As explained in Section 3, since the results obtained from the prescriptive model do not match those obtained with simulations, we need to modify the model in order to make them consistent. Our conjecture is that the discrepancy in performance between the prescriptive model and the simulated robot swarm is due to the fact that, as the number of robots grows, interference between robots reduces p_{ca} . This is because the robots spend time avoiding collisions and because the robots stopping in the black areas prevent other robots from accessing them. These aspects are not considered explicitly in the model. Reducing p_{ca} in the model allows us to obtain results that are closer to those obtained in simulation. For 10 robots there is no need to modify p_{ca} , as the results already match. For 20 robots and $p_{ca} = 0.05$, we observe that Property 1 is satisfied: robots form an aggregate in less than 1,000 seconds with probability 0.53. This matches the results obtained in simulation. For 50 robots we set $p_{ca} = 0.04$, which gives a probability of 0.01. Table III presents a comparison between the number of successful aggregates obtained before 1,000 seconds obtained in simulation and those obtained with model checking with the old and p_{ca} new values.

To test Property 2, we perform 100 runs of the simulated experiments for 10,000 seconds with the three group sizes. In the experiments, we measure whether the robot swarm satisfies Property 2, that is, whether a complete aggregate, once formed, lasts more than 10 seconds. In all the cases in which a complete aggregate was formed before 10,000 seconds, Property 2 was satisfied.

Videos of the simulated experiments are available in the supplementary material [Brambilla et al. 2014].

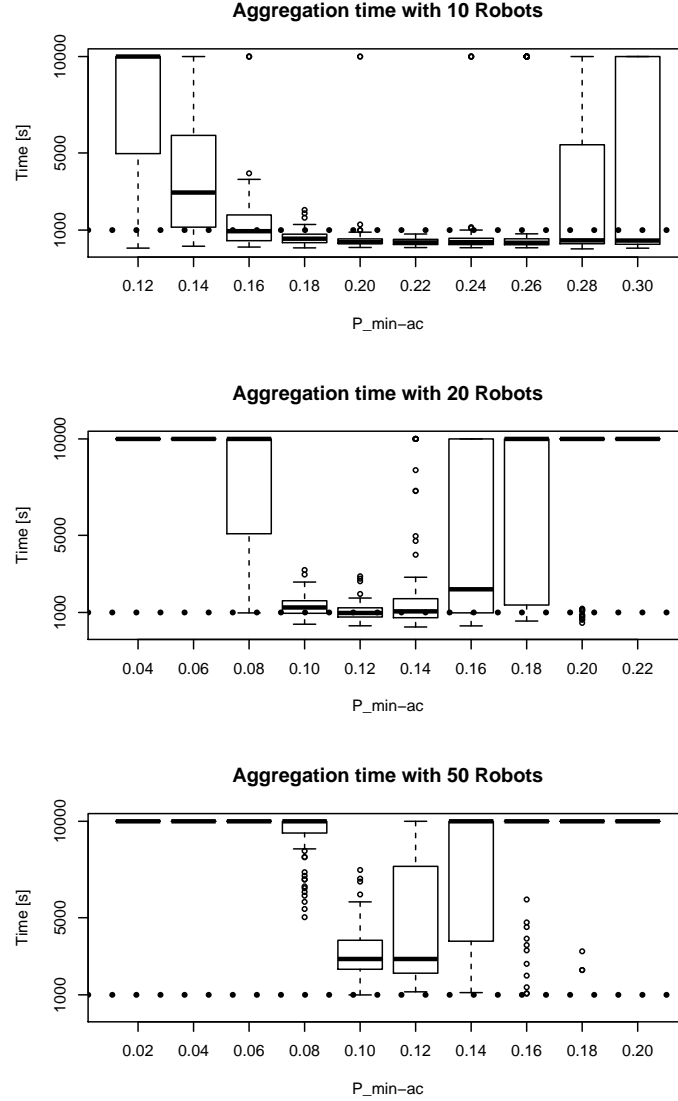


Fig. 4. The results obtained with the ARGoS simulator. The graphs show the time at which the experiment is stopped. This time is less than 10,000 seconds in case the aggregate is formed, or equal to 10,000 seconds in case the aggregate is not formed. Results are presented for different p_{min-ac} over 100 runs for 10, 20 and 50 robots.

Phase four: Robots – We perform 10 experiments with a group of 10 e-pucks in an arena identical to the simulated one. A screenshot of an experiment can be seen in Figure 5. Figure 6 shows a comparison between the time necessary for achieving aggregation obtained with the robots and in simulation. A video of a run is available in the supplementary material [Brambilla et al. 2014].




Fig. 5. A picture of an experiment performed with 10 e-puck robots.




Fig. 6. A graph showing the empirical cumulative distribution $F_n(x)$ of the time necessary to achieve aggregation obtained with robots (10 runs) and in simulation (100 runs). In both cases $N_t = 10$ and $p_{min-ac} = 0.22$.

In 10 runs out of 10, both Property 1 and Property 2 were satisfied. The results obtained with the robots are in line with those obtained in simulation.

The obtained robot swarms is able to aggregate satisfying the required properties. For this reason, there is no need to further update the prescriptive model and we can declare the process completed.

4.2. Foraging

In the second case study, we tackle foraging. In the simplest form of foraging, robots harvest *objects* and store them in the *nest*. The objects can be scattered in random positions or located in specific areas in the environment called *sources*. Foraging can be seen as an abstraction of more complex and realistic applications, such as search and rescue, land mine removal, waste cleaning and automated warehouse operation.

The number of objects retrieved typically depends on the number of robots: a single robot can perform foraging alone, but additional robots could be added to increase the performance of the swarm as robots working in parallel are able to retrieve more objects per time unit than a single robots. However, when the density of the robots in the environment increases, the performance of each single robot may be reduced due to interference [Lerman and Galstyan 2002; Pini et al. 2009].

In this case study, we assume that the robotic platform is given: the task must be tackled using e-pucks [Mondada et al. 2009]. The e-puck does not have the manipulation capabilities to interact with physical objects, so we consider an abstract version of foraging: instead of interacting with objects, e-pucks interact with TAM devices [Brutschy et al. 2010]. The TAM is a device similar to a booth, in which a robot can enter. It has a system of light barriers to sense the presence of a robot, and an LED that can be used to communicate information about its internal state. In this case study, TAMs are used to simulate the manipulation of objects: an e-puck can enter in a TAM, wait a fixed time and leave to simulate harvesting or storing an object.

The arena comprises 20 TAMs: 5 TAMs on the north wall act as the nest, each of these TAMs is a storing location; 15 TAMs on the other walls act as sources, locations where objects can appear. At any given time, in the arena there are O objects available, that is, a new object appears as soon as one is harvested by a robot. We perform experiments in which O equals $\{2, 4, 6, 8, 10\}$. The number of available storing locations depends on the number of robots currently storing an object: it can vary from 5, when no robot is using a storing location, to 0 if all are in use.

The state of a TAM is encoded using colors: green when the TAM is available for storage; blue when the TAM has an object available for harvesting; red when the TAM is busy, that is, a TAM in which a robot is currently harvesting or storing an object; off/black when the TAM is unavailable.

The environment is enclosed in $2\text{ m} \times 2\text{ m}$ square (see Figure 7 and Figure 12). Note that there is no globally perceivable clue in the environment that informs the robots on the position of the nest, differently from many other foraging studies (See Brambilla et al. [2013] for a review including work on foraging).

To allow robots to see the TAMs, we use e-pucks equipped with an omnidirectional camera.² Using the omnidirectional camera, robots can see the LEDs of the TAMs within a range of 0.5 m.

In the following, we will apply the 4-phase process presented in Section 3. In the foraging case study, we use the continuous time version of the Markov chain model, to model more easily the durations of some actions, such as harvesting and storing an object.

Phase one: Properties – In foraging, the main requirement is that the swarm retrieves at least a certain number of objects within a fixed time:

$$R\text{"obj_ret"} \geq k \quad [C \leq 600] \quad (3)$$

where $R\text{"obj_ret"} \geq k$ indicates that we are interested that the expected value of the reward "obj_ret" is greater or equal than k , and $C \leq 600$ indicates that we are inter-

²See <http://www.gettronic.com> for more details.




Fig. 7. A screenshot of the simulated version of foraging using 20 robots. Green lighted TAMs signal storage locations, blue lighted TAMs signal objects to be taken, dark TAMs are not available.

Table IV. The value of k necessary to satisfy Property 3 for different values of N_t , the number of robots composing the swarm, and values of O , the number of objects available in the environment at any given time.

N_t	O	k	N_t	O	k
20	2	45	10	6	40
20	4	55	20	6	65
20	6	65	50	6	90
20	8	75	100	6	75
20	10	85			

ested in the cumulative value over 600 seconds. The number k of objects that we wish to retrieve depends on N_t , the number of robots composing the swarm, and on O , the number of objects available in the environment at any given time; see Table IV.

Another requirement is on the *worst case performance*, that is, we want to ensure that the robot swarm is able to retrieve at least a minimum number of objects in 600 seconds. Model checking allows us to formally verify this condition since we can compute not only the expected value, but also its cumulative distribution (or, conversely, the density function). Formally, the second requirement is defined as:

$$P > 0.90 \ [\ F \leq 600 \ (\text{obj_ret} > 40) \] \quad (4)$$

In natural language, Property 4 can be expressed as: is it true with probability greater than 0.90 ($P > 0.90$) that at least 40 objects are retrieved ($\text{obj_ret} > 40$) in less than 600 seconds ($F \leq 600$)? To simplify the discussion, we verify Property 4 only in the case where $N_t = 20$ and $O = 6$.

Phase two: Model – To build the prescriptive model, we consider the different actions that a robot must perform. We then associate a state of the Markov chain to each of these actions.

A robot searches for objects by performing a random walk in the environment (So state). Once an object is found, the robot tries to harvest it (H state); in case of multiple objects in range, the robot goes towards the closest one. If the harvest action is unsuccessful, because, for instance, another robot harvests the object, the robot goes back to searching. When the object is reached, the robot waits inside the TAM for a fixed amount of time until the object is harvested (Hw state). Once the robot has harvested an object, it proceeds to search for the nest by performing random walk (Sn state). As soon as an available storage location is found, the robot tries to store the carried object (ST state); also in this case, the closest storage location is approached if multiple storing locations are seen. If the store action is unsuccessful, the robot searches for another storage location until the object is stored. Similar to the harvest operation, also in this case the robot waits inside a TAM for a fixed amount of time until the object is stored (STw state). A successful store operation increases the object counter (obj_ret). The robot then searches for a new object to harvest.

Robots always try to avoid collisions with obstacles and other robots. Practically, this produces two behaviors: when a robot is trying to enter a TAM (state H or state ST), it follows a vector that is the sum of a vector pointing to the desired destination and a vector pointing away from the closest obstacle. When a robot is performing random walk without a specific destination instead (state So or state Sn), if it encounters an obstacle or another robot, it starts turning on the spot for a random number of steps and then it begins again to move straight. This random number of steps follows a geometrical distribution. We model these different reactions in two different ways: in the first case, the action of the robot is not significantly disturbed, as the robot performs only a slight change of trajectory towards its goal. For this reason, this first kind of collision avoidance affects only the time to complete the action, but does not change the behavior of the robot. In the second case, instead, the robot completely changes its direction to avoid a collision, resulting in a significant change in its behavior and its chance to find objects or the nest. For this reason, the second kind of collision avoidance is modeled by adding two states: state Ao in case the robot is avoiding a collision when searching for an object, and state An in case the robot is avoiding a collision when searching for the nest.

See Figure 8 for a complete view of the prescriptive model.

We now have the structure of the behavior that the robots should follow. We need to assign values to the transition rates. We compute the transition rates considering the behavior of a single robot. For the macroscopic model, the rates are then multiplied by the current number of robots in the related state, as illustrated in Figure 8.

All the rates involved in the definition of the model depend on the geometrical characteristics of the environment and/or on the behavior of the robots. Unfortunately, differently from the previous case study, we cannot completely define them a priori since it is impossible to identify the correct value of the parameters involved in their definition without experimental data. Note that the goal of this phase is not to create a model that is as precise as possible, but one that can be used by us to develop and improve the desired robot swarm. Since we do not have experimental data, the model we are creating is largely arbitrary. Other valid choices could have been made. We make some working hypotheses about the system that can be subject to refinements or changes in the subsequent phases, should they prove not to be sufficiently accurate or correct. In particular, parameters will be fitted once experimental data are available, that is, in phase three.

In the following, we present how each rate is defined. We define $\lambda_{So \rightarrow H}$, the rate at which a robot finds an object, as proportional to the density of available objects in the




Fig. 8. The continuous time Markov chain used to model foraging. H is the *harvest* state; Hw is the *wait to harvest* state; So is the *search object* state; Ao is the *avoid (while searching for an object)* state; ST is the *store* state; STw is the *wait to store* state; Sn is the *Search nest* state; An is the *Avoid (while searching for the nest)* state; Transitions are labeled with their respective rates multiplied by the number of robots currently in that state: $\lambda_{i \rightarrow j}$ is the rate at which an individual robot moves from state i to state j ; N_i is the number of robots currently in state i . To compute the expected number of objects retrieved, we keep track of the number of times the transition from STw to So, labeled **Obj_ret**, happens.

environment:

$$\lambda_{So \rightarrow H} = \alpha \frac{O}{A},$$

where O is the number of objects available at any given time, A is the area of the environment and α is a parameter.

We define $\lambda_{So \rightarrow Ao}$, the rate at which a robot searching for an object finds another robot and then performs obstacle avoidance, as proportional to the density of robots in the environment:

$$\lambda_{So \rightarrow Ao} = \beta \frac{N_r}{A},$$

where N_r is the total number of robots in the environment, and β is a parameter.

Rates $\lambda_{Sn \rightarrow ST}$ and $\lambda_{Sn \rightarrow An}$ are defined similarly, considering the number of storage locations instead of the number of objects:

$$\lambda_{Sn \rightarrow ST} = \gamma \frac{D}{A},$$

$$\lambda_{Sn \rightarrow An} = \delta \frac{N_r}{A},$$

where D is the number of storage locations and γ and δ are two parameters.

Rates $\lambda_{Ao \rightarrow So}$ and $\lambda_{An \rightarrow Sn}$ depend on the time necessary for a robot to perform obstacle avoidance. As said before, if a robot encounters an obstacle or another robot while searching for objects or for the nest, it performs collision avoidance by turning on the spot for a random number of steps distributed geometrically and then it starts again searching for objects or the nest. The rate at which a robot moves from collision avoidance back to search is thus:

$$\lambda_{Ao \rightarrow So} = \lambda_{An \rightarrow Sn} = p_{oa},$$

where p_{oa} is the parameter of the geometrical distribution.

We define $\lambda_{H \rightarrow Hw}$ and $\lambda_{ST \rightarrow STw}$, the rates at which a robot going towards an object-TAM or a storage-TAM manages to enter it, as the reciprocal of the time necessary to get in

the TAM, counted from the instant in which the robot sees it:

$$\lambda_{H \rightarrow Hw} = \lambda_{ST \rightarrow STw} = \left(\frac{r}{s}\right)^{-1},$$

where r is the range at which a robot sees a TAM and s is the forward speed of a robot.

A robot trying to enter a TAM is not always successful, other robots may “steal” its object by occupying the storage location before it can do it. This means that not all robots going towards a TAM enter it. Some are interrupted by other robots and thus are forced to search for another available TAM. This is modeled by the transition $H \rightarrow So$ and $ST \rightarrow Sn$. We define $\lambda_{H \rightarrow So}$ and $\lambda_{ST \rightarrow Sn}$, the related rates, as proportional to the density of robots in the environment:

$$\lambda_{H \rightarrow So} = \epsilon \frac{N_r}{A},$$

$$\lambda_{ST \rightarrow Sn} = \eta \frac{N_r}{A},$$

where ϵ and η are parameters. We expect ϵ and η to have different values, as storage locations are all next to each other, generating more interference, while objects to harvest in general are evenly distributed along the walls of the environment.

The last rates we need to define are $\lambda_{Hw \rightarrow Sn}$ and $\lambda_{STw \rightarrow So}$, the rates at which robots in the TAM complete their operations and exit. These rates are the reciprocal of the time spent by a robot in a TAM:

$$\lambda_{Hw \rightarrow Sn} = \lambda_{STw \rightarrow So} = (t_{TAM})^{-1},$$

where t_{TAM} is the time spent by a robot in a TAM.

Since we do not have empirical data to estimate the parameters, at this point we cannot use model checking to compute the expected number of objects retrieved, or whether the desired properties are satisfied. However, even without empirical data, we can use the model to improve the behavior of the robots. For example, by analyzing the model, we can observe that increasing the rate at which the robots find objects or storage locations—that is, increasing $\lambda_{So \rightarrow H}$ and $\lambda_{Sn \rightarrow ST}$ —results in an increase of objects retrieved.

In order to increase these rates, we cannot modify the number of objects available at any given time or the number of storage locations, since they are given. We could act on the parameters, but it is not clear how to change the behavior of the robots to increase these parameters. Even though we cannot change the dimensions of the environment, we can change the size of the area effectively covered by the robots. In other terms, we can change the behavior of the robots so that they do not cover the whole environment when searching for objects or storage locations. In particular, we could let robots avoid places where they know they will not find anything useful.

In the behavior defined before, robots searching for objects and for the nest go straight until they find an obstacle such as a wall or another robot. This means that robots that are carrying an object while searching for the nest may go close to other available objects, interfering with robots not carrying objects. Similarly, robots searching for objects often go close to the storing locations, interfering with the other robots. A possible solution is that robots searching for objects avoid storing locations as soon as they see them and, similarly, robots searching for the nest avoid objects as soon as they see them. This improved behavior is depicted in Figure 9.




Fig. 9. The modification of the behavior used to reduce interference and reduce the area searched by the robot. A robot (depicted as a green circle with a yellow arrow on top) carrying an object performs collision avoidance (full arrow) as soon as it sees an other object, instead of reaching to the wall (dotted arrow). The light blue circular sections represent the areas in which a robot sees an object. The same applies to robot searching for the nest, even though it is not displayed in the figure.

This improvement in the behavior can be modeled by decreasing the value of A . The rates are updated in the following way:

$$\lambda_{\text{So} \rightarrow \text{H}} = \alpha \frac{O}{A_o},$$

where A_o is the area explored by the robots searching for objects and avoiding storage locations, with $A_o < A$;

$$\lambda_{\text{Sn} \rightarrow \text{ST}} = \gamma \frac{D}{A_n},$$

where A_n is the area explored by the robots searching for storage locations and avoiding objects, with $A_n < A$; All other rates are left unchanged since the density of the robots involved in their definition does not change. For example, consider $\lambda_{\text{So} \rightarrow \text{Ao}}$. The area involved in the definition of this rate is reduced, as explained above. However, also the number of robots operating in that area is reduced. In other words, even though the area considered is reduced, the density of robots does not change.

More complex improvements, such as task allocation mechanisms, could be implemented to further increase the performance of the system, should the obtained performance not be sufficient. However, for the sake of brevity and clarity we limit our design process to the simple improvement presented above.

Phase three: Simulation – In this phase, we implement the foraging robot swarm using the ARGoS simulator [Pincioli et al. 2012].

The prescriptive model developed in phase two provides us with a detailed blueprint to implement the robot swarm: the behavior of the individual robot can be implemented using a finite state machine that resembles the Markov chain defined in phase two. Nonetheless, some implementation details, such as how robots stop inside a TAM, have been ignored in the prescriptive model, in order to focus on the more important details at design time and have to be programmed explicitly at this moment.

We measured the number of objects retrieved over 600 seconds on 100 runs. We first performed experiments using the initial behavior and then using the improved one, as

Table V. The numerical parameters used in the foraging prescriptive model. The values are obtained from the experimental data obtained in phase three.

	Value		Value
O	$\{2, 4, 6, 8, 10\}$	p_{oa}	0.20
N_r	$\{10, 20, 50, 100\}$	r	0.5 m
D	5	s	0.1 m s^{-1}
l	2 m	ϵ	6.5×10^{-2}
A	4 m^2	η	1.01×10^{-1}
α	4×10^{-2}	t_{TAM}	3 s
β	1×10^{-2}	A_o	3 m^2
γ	3.9×10^{-2}	A_n	3.7 m^2
δ	9×10^{-3}		

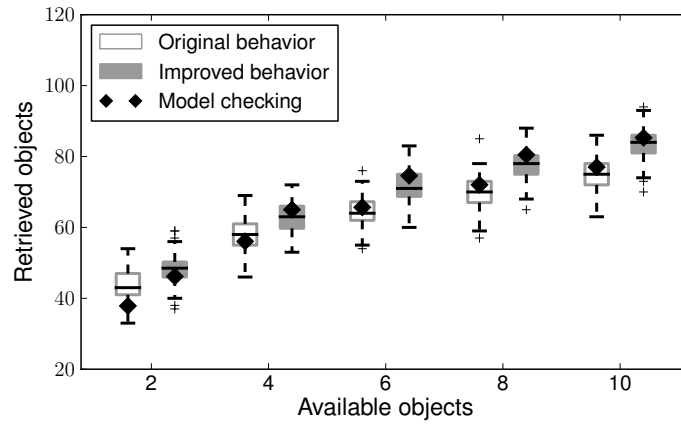


Fig. 10. A comparison between the results obtained using 20 robots with the original behavior and with the improved one for different values of O , which is the number of objects available at any time. Box plots show results obtained over 100 experimental runs using the ARGoS simulator, while diamonds show the expected results obtained with PRISM.

explained in the previous phase. Table V shows the parameter used for model checking derived from the experimental data.

We can now compute the expected number of objects retrieved in 600 seconds using model checking on the developed model and compare these values with the results obtained from the simulated experiments.

Figure 10 shows the results obtained in simulation together with the expected results predicted by the prescriptive model. These results have been obtained using 20 robots with different values of O , the number of objects available at any time. Figure 11 shows the results obtained with $O = 6$ and a different number of robots. Table VI shows whether Property 3 is satisfied.

From Figure 10 it is possible to observe that indeed the behavior improvement introduced in phase two significantly increases the number of objects retrieved. The improved behavior is always significantly better than its counterpart—Wilcoxon test with $p < 0.01$.

The correspondence between the results obtained from the prescriptive model and the ones obtained from the simulations is quite good, even though not perfect. For our goals and purposes, the model captures qualitatively the behavior of the robot swarm, thus it is not necessary to further refine it.

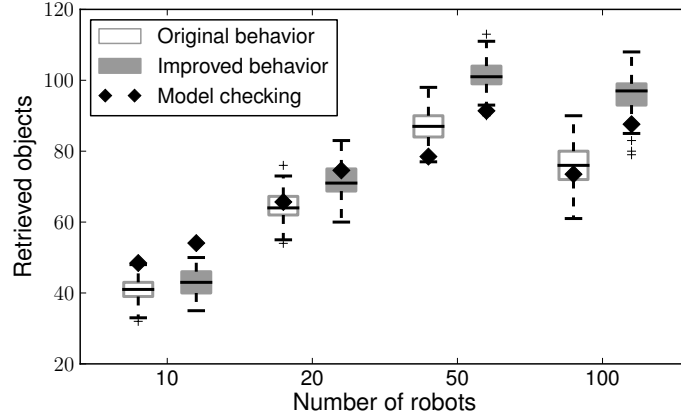


Fig. 11. The number of object retrieved with different swarm sizes and $O = 6$. Box plots show results obtained over 100 experimental runs using the ARGoS simulator, while diamonds show the expected results obtained with PRISM.

Table VI. A table presenting whether the simulated systems are able to satisfy Property 3. k is the threshold on the expected number of objects retrieved. Property 3 is satisfied if the values obtained in simulations are greater or equal than k . The results are presented both for the original and the improved behavior, showing whether they satisfy Property 3 and the value of the median.

N_t	O	k	Original	Improved	N_t	O	k	Original	Improved
20	2	45	X (43)	✓ (48)	10	6	40	✓ (41)	✓ (43)
20	4	55	✓ (58)	✓ (63)	20	6	65	X (64)	✓ (71)
20	6	65	X (64)	✓ (71)	50	6	90	X (87)	✓ (101)
20	8	75	X (70)	✓ (78)	100	6	75	✓ (76)	✓ (97)
20	10	85	X (75)	✓ (85)					

We also verify Property 4 using model checking and compare it with the results obtained from the simulations. Model checking tells us that Property 4 is not satisfied in the prescriptive model of the original behavior with 20 robots and $O = 2$. This matches the experimental results, where 15 runs over 100 resulted in less than 40 objects retrieved. Instead, with the improved behavior, property 4 is satisfied. This matches the experimental results, where only 2 runs over 100 resulted in less than 40 objects retrieved.

All experimental data can be found in the supplementary material [Brambilla et al. 2014].

Phase four: Robots – We performed 10 experiments with a group of 20 e-pucks in an arena identical to the simulated one. A picture of an experiment can be seen in Figure 12. Videos of the performed experiments can be found in the supplementary material [Brambilla et al. 2014]. Figure 13 shows that the results obtained with real robots and simulated robots are very similar. Property 4 is satisfied. There is no need to update the model and thus we can declare the process completed.




Fig. 12. A picture of an experiment performed with 20 e-puck robots and $O = 6$.




Fig. 13. A graph showing the empirical cumulative distribution $F_n(x)$ of the number of object retrieved using robots (10 runs) and in simulation (100 runs). In both cases $N_t = 20$ $O = 6$.

4.3. Discussion

The two case studies presented in the paper show that using property-driven design we were able to develop two robot swarms that tackle successfully aggregation and foraging: all the required properties are satisfied.

As shown, with property-driven design it is possible to analyze and develop behaviors characterized both by numerical and non-numerical parameters. For example, in the aggregation case we analyzed the effects of changing the probability to leave a black

area, whereas in the foraging case we analyzed the effects of changing the exploration behavior.

One of the main advantages of property-driven design is that it allows the developer to focus on the design of the system rather than on its implementation: by focusing the designing efforts on the abstract model of the system, the developer can concentrate on the important aspects of the system to create because “whereas a simulation should include as much detail as possible, a good model should include as little as possible” [Smith 1978]. As an example, in the foraging case study we leveraged the prescriptive model developed in phase two of property-driven design to identify possible improvements of the partially realized system: we were able to identify the important aspects of the system being designed, which allowed us to avoid wasting time on improving other non-relevant aspects.

Another advantage of property-driven design is the reduced risk of developing a robot swarm that does not satisfy the requirements. Up to now, there was no clear way to specify the requirements of a robot swarm. In property-driven design, requirements are specified in a formal way at the beginning of the development process in terms of desired properties. Moreover, thanks to model checking it is possible to evaluate whether the robot swarm fulfills such properties at each step of the design and development process. This advantage of property-driven design has been highlighted by both case studies.

Finally, property-driven design addresses also the problem of the low re-usability of solutions in swarm robotics. Usually behaviors for robot swarms are developed in a disposable way. This is due to the fact that there is no clear distinction between the design and the implementation. Thus, if a different hardware platform is available, or a slightly different task is tackled, it is necessary to start from scratch. With property-driven design instead, the prescriptive model developed in phase two can be partially or completely reused: i) the model is hardware independent, so that it can be adapted to the available robots, or even guide the process of deciding the best robot to use; and ii) the model can be extended to deal with new properties and verify if they are satisfied even without testing the system in simulation or with robots. The reusability of the prescriptive model reduces the risk that designers “reinvent the wheel” each time they develop a robot swarm. For example, the model of foraging developed in the presented case study could be easily adapted for robots with more sophisticated manipulation capabilities. In the future, it is also possible to imagine a set of publicly available models for swarm robotics applications that can be reused and modified by other developers.

The development process of the case studies presented in the paper highlights also some issues with property-driven design.

The main issue is that ultimately the step from the prescriptive model to its implementation remains in the hands of the developer. Nonetheless, the prescriptive model can be used as a blueprint for the implementation process, providing the developer with a valuable tool to obtain robot swarms with provable properties.

Another issue is the strong reliance on modeling. Modeling robot swarms is a difficult task on its own: robot-to-robot interactions, spatial and temporal features and interference are difficult to completely describe using models. Luckily, modeling robot swarms has been the focus of a large number of studies (see two reviews of the literature by Brambilla et al. [2013] and Lerman et al. [2005]) providing a solid theoretical foundation to property-driven design.

5. CONCLUSIONS

Property-driven design is a top-down design method based on prescriptive modeling and model checking: the desired robot swarm is first described using a set of properties;

subsequently a prescriptive model of the robot swarm is created; the prescriptive model is used as a blueprint for the implementation of the robot swarm first in simulation and then with robots.

Property-driven design is conceived to be part of swarm engineering: the systematic application of scientific and technical knowledge to specify requirements, design, realize, verify, validate, operate and maintain a swarm intelligence system. Up to now, the design and development of a robot swarm is performed using a code-and-fix approach based completely on the ingenuity and experience of the developer who does not have any scientific or technical support in his activity. Property-driven design aims at providing such scientific and technical support, with many advantages compared to the traditional unstructured approach.

In this paper, we demonstrated, by tackling two different case studies, that property-driven design is an effective method for the design and development of robot swarms. In the future we plan to apply property-driven design to different and more complex tasks, possibly using different modeling approaches. Also, we aim at integrating this design method with automatic design approaches, such as evolutionary robotics or automatic modular design [Francesca et al. 2014], in order to obtain automatically robot swarms that satisfy some desired properties. This could provide a solution to the open problem of deriving the individual behaviors from a swarm-level model.

APPENDIX

Property-driven design is based on prescriptive modeling and model checking. Model checking requires two components: a model of the system to check and a set of properties that the system must satisfy. Among the several possible languages and tools, we chose Markov chains and probabilistic temporal logics. Markov chains and probabilistic temporal logics have been proved to be well suited for model checking in swarm robotics [Konur et al. 2012; Brambilla et al. 2012] due to their simplicity and expressive power.

In this section, we give a short introduction on Markov chains, probabilistic temporal logics and model checking.

I. The model: Markov chains

A common way to model swarm robotics systems is through the use of Markov chains [Lerman et al. 2005]. Markov chains are used to model the behavior of the robots: *states* might represent an action that a robot performs, such as random walk or grasp object, or might represent an area in which a robot is located, for example, in the nest. Transitions link two states and are activated through *transition conditions* such as obstacle seen and object grasped.

Markov chains can be used to model a swarm robotics system in two ways: as a *microscopic Markov model*, that is, a model that considers each individual robot, and as a *macroscopic Markov model*, that is, a model that consider the swarm as a whole.

A microscopic Markov model describes the behavior of the individual robots and their interactions. In a microscopic Markov model, the collective behavior of the swarm is the *and*-composition of the individual Markov chains.

A macroscopic Markov model describes the swarm as a whole, without considering the individual robots composing it. In general, as explained by Lerman et al. [2005], a macroscopic Markov model is composed of an augmented Markov chain: the Markov chain describing the behavior of a generic individual of the swarm is augmented by associating a counter to each state. Counters are used to keep track of the number of robots that are in the associated state in any given moment. This is different from rate equations in which only the proportion of the robots is tracked. Examples of macroscopic Markov models can be found in Section 4, in particular in Figure 2 and Figure 8.

When compared to macroscopic Markov models, microscopic Markov models give a finer description of the robots and their interactions. However, in a microscopic Markov model the number of states grows exponential with the number of robots, making microscopic Markov models computationally intractable in the majority of cases. Since microscopic Markov models




Fig. 14. A simple Markov chain (on the left) and part of its computation tree (on the right).

are difficult to analyze, the great majority of models in swarm robotics are macroscopic [Brambilla et al. 2013; Lerman et al. 2005].

Time in Markov chains can be modeled in two different ways: discrete and continuous. In discrete time Markov chains (DTMC), time assumes only values in \mathbb{Z}^+ , whereas in continuous time Markov chains (CTMC), time can assume any value in \mathbb{R}^+ . Practically, one of the main difference between DTMC and CTMC is in what transition parameters represent. In DTMC, transition parameters represent the probability p of moving from a state to another; for this reason, transition parameters for DTMC must be in the interval $[0, 1]$. In CTMC, instead, transitions parameters represents the rate λ at which the event of moving from a state to another happens, that is, the time taken to move from a state to another follows an exponential distribution of parameter $\lambda \in (0, \infty)$.

The choice on how to model time depends on the system to describe: in case time is not the main aspect and can be easily discretized, DTMC are more convenient; on the contrary, when it is important to keep precisely track of times, DTMC should be preferred. Note that, in general, it is possible to model the same system using a DTMC or a CTMC without losing expressive power [Serfozo 1979].

II. The properties: probabilistic temporal logics

The most common way to formally express properties for model checking is through the use of logic predicates. Among the many formal logic systems, we consider *probabilistic computation tree logic* (PCTL)³. PCTL is well suited for swarm robotics systems as it can capture both time-related and stochastic aspects.

PCTL [Hansson and Jonsson 1994] is based on the concept of computation tree, a potentially infinite rooted tree in which the root is the initial state of a corresponding Markov chain, and each node is a possible state of the system. Edges link a state with its next possible states. Each path on the tree represents a possible execution of the system. Since a sequence of nodes represents the time evolution of a system, the transition from a node to a following one is usually called a time-step. In discrete time Markov chains, this time-step is fixed, while in continuous time Markov chains, this time-step is exponentially distributed with parameter depending on the current state. An example of a simple Markov chain and its computation tree is displayed in Figure 14.

A computation tree can be used to express temporal properties, such as *eventually the system will reach state X* or *if the system starts from state α then it will never reach state β* . Such properties can be expressed using computation tree logic (CTL).

PCTL extends CTL by introducing probabilities. It is thus possible to express properties such as *α will eventually become true with probability 0.45* or *there is a 0.7 probability that α will hold true for 10 seconds*. In this paper, we do not use the formal syntax of PCTL. We instead use the syntax of PRISM: the tool used for model checking. A formal introduction to the syntax of PCTL can be found in the work by Ciesinski and Größer [2004].

³Note that PCTL can be used only with discrete time Markov chains. For continuous time Markov chains, it is necessary to use *continuous stochastic logic* (CSL). For our goals and purposes however, the two logics are equivalent. Thus, for the sake of simplicity, we refer to PCTL also when dealing with CTMCs, even though this is formally incorrect.

Probabilistic and temporal operators make PCTL a very flexible and powerful logic, which can be used to express many interesting properties of particular interest for swarm robotics systems.

III. Model checking: complete and statistical model checking

Having defined a model and a set of properties, we now have all the elements to perform model checking.

Model checking can be used in safe-critical applications in which simulations and experiments might not be enough to guarantee the correctness of a system. In fact, simulations and experiments can only test a subset of all possible execution scenarios of a system. Model checking, instead, formally verifies that a property holds true for all possible executions of a system.

Model checking has a number of advantages compared to more traditional ways to analyze models of swarm robotics systems, such as fluid flow analysis [Zarzhitsky et al. 2005].

Not only model checking allows the user to verify that a model satisfy a specific probabilistic property (as, for example, $P_{\geq 0.75}[\phi] \text{ ? TRUE}$), but also to obtain quantitative results, that is, to compute with which probability the model satisfies it ($P_{\tau}[\phi] \text{ ? } 0.86$). This characteristic is very useful to find the best parameters of a model that maximize the probability to satisfy a specific property. Moreover, it is possible to augment a Markov chain using *rewards*, real valued quantities that can be assigned to states or transitions. Using model checking it is possible to compute not only the steady state value of these rewards, but also their probability distribution. This would be impossible with analysis based on rate equations, as they only provide the expected value of the observed variable. Model checking can also be used to produce counter examples: traces of execution of a system in which a property is not satisfied. Additionally, the use of PCTL allows the developer to express properties that are difficult or even impossible to express using algebraic mathematics.

A limit of model checking is that, in general, it is computationally unfeasible to analyze models composed of a high number of states. State-of-the-art model checkers cannot handle models larger than 10^{10} states [Kwiatkowska et al. 2004]. A way to overcome this problem is *statistical model checking*. Statistical model checking, also known as approximate model checking, is a novel approach to model checking [Nimal 2010]. Compared to traditional model checking, statistical model checking does not explore completely the state space of a model. Instead, it samples a large but limited number of executions of the model and uses statistic estimators to compute the result.

Using statistical model checking it is thus possible to perform model checking also on very large models, such as microscopic models of swarm robotics systems. In Massink et al. [2013], we have applied statistical model checking to a model of a swarm robotics system and showed that the obtained results were consistent with those obtained using other approaches, such as physics-based simulation, Monte Carlo simulation and ordinary differential equations.

ACKNOWLEDGMENTS

The authors would like to thank Arne Brutschy for his support with the TAM.

REFERENCES

- W. Banzhaf and N. Pillay. 2007. Why complex systems engineering needs biological development. *Complexity* 13, 2 (2007), 12–21.
- K. Beck. 2003. *Test-driven Development: By Example*. Addison-Wesley, Boston, MA.
- S. Berman, A. Halasz, M.A. Hsieh, and V. Kumar. 2009. Optimized stochastic policies for task allocation in swarms of robots. *IEEE Transactions on Robotics* 25, 4 (2009), 927–937.
- S. Berman, V. Kumar, and R. Nagpal. 2011. Design of control policies for spatially inhomogeneous robot swarms with application to commercial pollination. In *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE press, 378–385.
- R.H. Bordini. 2009. *Multi-Agent Programming: Languages, Tools and Applications*. Vol. 2. Springer-Verlag, New York, NY.

- M. Brambilla, M. Dorigo, and M. Birattari. 2014. Property-driven design for robot swarms: Supplementary material. (2014). Supplementary material available at https://dl.dropboxusercontent.com/u/33171834/property-driven_supp_mat.zip. The material is distributed via Dropbox to ensure anonymity of consultation. It will be posted online on the IRIDIA supplementary pages, should the paper be accepted.
- M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo. 2013. Swarm robotics: A review from the swarm engineering perspective. *Swarm Intelligence* 7, 1 (2013), 1–41.
- M. Brambilla, C. Pinciroli, M. Birattari, and M. Dorigo. 2012. Property-driven design for swarm robotics. In *Proceedings of the AAMAS 2012*. IFAAMAS, 139–146.
- A. Brutschy, G. Pini, N. Baiboun, A. Decugnière, and M. Birattari. 2010. *The IRIDIA-TAM: A device for task abstraction for the e-puck robot*. Technical Report TR/IRIDIA/2010-015. IRIDIA, ULB.
- F. Ciesinski and M. Größer. 2004. On probabilistic computation tree logic. In *Validation of Stochastic Systems*. Lecture Notes in Computer Science, Vol. 2925. Springer, Berlin, Heidelberg, 333–355.
- C. Dixon, A. Winfield, and M. Fisher. 2012. Towards temporal verification of swarm robotic systems. *Robotics and Autonomous Systems* 60, 11 (2012), 1429–1441.
- M. Dorigo, M. Birattari, and M. Brambilla. 2014. Swarm robotics. *Scholarpedia* 9, 1 (2014), 1463.
- G. Francesca, M. Brambilla, A. Brutschy, V. Trianni, and M. Birattari. 2014. *AutoMoDe: A novel approach to the automatic design of control software for robot swarms*. Technical Report TR/IRIDIA/2014-001. IRIDIA, Université Libre de Bruxelles, Brussels, Belgium.
- D. Goldberg and M. J. Mataric. 2001. Design and evaluation of robust behavior-based controllers for distributed multi-Robot collection tasks. In *Robot Teams: From Diversity to Polymorphism*. A K Peters Ltd, 315–344.
- A. Gutiérrez, A. Campo, M. Dorigo, J. Donate, F. Monasterio-Huelin, and L. Magdalena. 2009. Open e-puck range & bearing miniaturized board for local communication in swarm robotics. In *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE Press, Piscataway, NJ, 3111–3116.
- H. Hamann and H. Wörn. 2008. A framework of space–time continuous models for algorithm design in swarm robotics. *Swarm Intelligence* 2, 2 (2008), 209–239.
- H. Hansson and B. Jonsson. 1994. A logic for reasoning about time and reliability. *Formal Aspects of Computing* 6, 5 (1994), 512–535.
- R. Jeanson, C. Rivault, J.-L. Deneubourg, S. Blanco, R. Fournier, C. Jost, and G. Theraulaz. 2005. Self-organized aggregation in cockroaches. *Animal Behaviour* 69, 1 (2005), 169–180.
- S. Kazadi, J. R. Lee, and J. Lee. 2009. Model independence in swarm robotics. *International Journal of Intelligent Computing and Cybernetics, Special Issue on Swarm Robotics* 2, 4 (2009), 672–694.
- S. Konur, C. Dixon, and M. Fisher. 2012. Analysing robot swarm behaviour via probabilistic model checking. *Robotics and Autonomous Systems* 60, 2 (2012), 199–213.
- M. Kwiatkowska, G. Norman, and D. Parker. 2004. Probabilistic symbolic model checking with PRISM: A hybrid approach. *International Journal on Software Tools for Technology Transfer* 6, 2 (2004), 128–142.
- K. Lerman and A. Galstyan. 2002. Mathematical model of foraging in a group of robots: Effect of interference. *Autonomous Robots* 13, 2 (2002), 127–141.
- K. Lerman, A. Martinoli, and A. Galstyan. 2005. A review of probabilistic macroscopic models for swarm robotic systems. In *Swarm robotics*. Lecture Notes in Computer Science, Vol. 3342. Springer, Berlin, Heidelberg, 143–152.
- M. Massink, M. Brambilla, D. Latella, M. Dorigo, and M. Birattari. 2013. On the use of Bio-PEPA for modelling and analysing collective behaviours in swarm robotics.

- Swarm Intelligence* 7, 2–3 (2013), 201–228.
- K. L. McMillan. 1993. *Symbolic Model Checking*. Springer, Berlin, Heidelberg.
- J. Miller and J. Mukerji. 2003. MDA guide version 1.0.1. <http://www.omg.org>. (2003).
- F. Mondada, M. Bonani, X. Raemy, J. Pugh, C. Cianci, A. Klapotcz, S. Magnenat, J.-C. Zufferey, D. Floreano, and A. Martinoli. 2009. The e-puck, a robot designed for education in engineering. In *Proceedings of the 9th Conference on Autonomous Robot Systems and Competitions*, Vol. 1. IPCB, 59–65.
- V. Nimal. 2010. *Statistical Approaches for Probabilistic Model Checking*. MSc Mini-project Dissertation. Oxford University Computing Laboratory.
- C. Pinciroli, V. Trianni, R. O’Grady, G. Pini, A. Brutschy, M. Brambilla, N. Mathews, E. Ferrante, G. Di Caro, F. Ducatelle, M. Birattari, L. M. Gambardella, and M. Dorigo. 2012. ARGoS: A modular, parallel, multi-engine simulator for multi-robot systems. *Swarm Intelligence* 6, 4 (2012), 271–295.
- G. Pini, A. Brutschy, M. Birattari, and M. Dorigo. 2009. Interference reduction through task partitioning in a robotic swarm. In *Sixth International Conference on Informatics in Control, Automation and Robotics (ICINCO)*. INSTICC Press, 52–59.
- E. Şahin. 2005. Swarm robotics: from sources of inspiration to domains of application. In *Swarm Robotics*. Lecture notes in computer science, Vol. 3342. Springer, Berlin, Heidelberg, 10–20.
- R. F. Serfozo. 1979. An equivalence between continuous and discrete time Markov decision processes. *Operations Research* 27, 3 (1979), 616–620.
- J. M. Smith. 1978. *Models in Ecology*. Cambridge University Press, Cambridge, MA.
- M. Wooldridge and N. R. Jennings. 1998. Pitfalls of agent-oriented development. In *Proceedings of the second international conference on Autonomous agents*. ACM Press, 385–391.
- F. Zambonelli, N. Jennings, and M. Wooldridge. 2001. Organisational abstractions for the analysis and design of multi-agent systems. In *Agent-Oriented Software Engineering*. Lecture Notes in Computer Science, Vol. 1957. Springer, Berlin, Heidelberg, 407–422.
- D. Zarzhitsky, D. Spears, D. Thayer, and W. Spears. 2005. Agent-based chemical plume tracing using fluid dynamics. In *Formal Approaches to Agent-Based Systems*. Lecture Notes in Computer Science, Vol. 3228. Springer, Berlin, Heidelberg, 146–160.

Received XXX; revised XXX; accepted XXX