

Comparing and Modeling Distributed Control Strategies for Miniature Self-Assembling Robots

William C. Evans, Grégory Mermoud and Alcherio Martinoli

Abstract—We propose two contrasting approaches to the scalable distributed control of a swarm of self-assembling miniaturized robots, specifically the formation of chains of a desired length: (1) a deterministic controller in which robots communicate with each other in order to directly limit the size of each chain, and (2) a probabilistic controller where the average chain size is controlled by the probability a robot will choose to leave its chain. We demonstrate the feasibility of both approaches by implementing them on a real swarm of Alice robots. Using Webots, a realistic simulator for mobile robotics, and macroscopic models based on the Chemical Reaction Network (CRN) framework, we investigate the limitations of the deterministic controller and demonstrate the existence of optimal parameters for the probabilistic controller where exploration and exploitation are well balanced, thus favoring the formation of larger chains.

I. INTRODUCTION

Aggregation is an ubiquitous phenomenon occurring at all scales: atoms and molecules [1], microfabricated parts [2], animals [3], and robotic systems, both with passive objects [4] and robots [5] as building blocks. Also, many functional structures found in nature are generated by a specific type of aggregation, self-assembly, which is essentially stochastic, reversible and structured [6], often yielding systems in which the stability of an aggregate depends on the relative positioning of its building blocks.

This paper investigates and models the self-assembly of miniaturized robots with minimal abilities in terms of sensing, actuation, and control. To this end, we use a swarm of Alice robots [7]. The Alice micro-robot is a 2cm cube, equipped with an extremely limited capacity for sensing and computation. In particular, the Alice has very poor odometry and no radio communication module (in its basic configuration), making deliberative motion planning an arduous and error-prone task. Therefore, in this paper we consider the self-assembly of Alice robots to be essentially stochastic, i.e. the robots perform a random walk and, upon collision, they can decide to aggregate. Local infrared communication allows aggregated robots to exchange some bits of information, albeit with poor reliability.

More specifically, we investigate the self-assembly of Alice robots into chains of a target size. We consider two approaches: (1) a deterministic approach where robots communicate with each other in order to determine the size

of each chain and adapt their behavior accordingly, and (2) a probabilistic approach where chain size is controlled by the probability that a robot will leave its current aggregate. Both controllers are described further in Section III-D.

We perform a detailed study of each controller using both Webots, a realistic simulator for mobile robotics, and macroscopic models based on Chemical Reaction Networks (CRN), a modeling framework that has been studied extensively in the context of biochemical systems [8]. This level of representation is very abstract, allowing one to capture the system as a whole, while still enabling the construction of exact and approximate numeric simulations of the system using the Gillespie method and its various optimizations [9], [10].

II. STATE OF THE ART

Recently, self-assembly and its unstructured equivalent, aggregation, have seen increasing popularity in distributed and modular robotics. This includes aggregation of passive objects mediated by mobile robots [4], [11], aggregation of Alice robots using a notably simple controller [12] and even mixed societies of robots and insects [13]. Among recent implementations of self-assembling robots at the macroscale, Gross et al. demonstrated self-assembly using the Swarmbot, a mobile robot equipped with a gripper [5]. However, because they rely on a full breadth of sensors, including cameras and a computationally intensive controller, their approach is not well suited to extreme miniaturization. Klavins et al. have demonstrated self-assembly of triangular robotic modules that slide passively on an air table [14]. Here, permanent magnets serve as a binding mechanism, and robots use explicit communication to execute a common graph grammar and determine their next course of action. Once again, we believe that this type of controller is a strongly limiting factor in the context of extreme miniaturization.

Previous attempts to use CRNs and biochemical modeling in robotics have been very successful [15]–[17]. Klavins and colleagues have used a similar framework for building a stochastic interpretation of their graph grammars [14]. The present paper is the first clear attempt to use the CRN toolbox, originally designed for the study of molecular systems, to model robotic self-assembly. Indeed, deterministic models of aggregation and flocking (which is conceptually similar to aggregation, but involves a coordinated motion of the aggregate) such as [18], [19] as well as graphical models of multi-robot systems such as [20] are interesting complements, from a system and control perspective, to our stochastic modeling approach, but do not take explicitly into

W. C. Evans, G. Mermoud and A. Martinoli are with the Distributed Intelligent Systems and Algorithms Laboratory, École Polytechnique Fédérale Lausanne, 1015 Lausanne, Vaud, Switzerland `firstname.lastname@epfl.ch`

This work was partially supported by SelfSys, a project sponsored by the Swiss research initiative Nano-Tera.ch.




Fig. 1. (a) The Alice mobile robot is a 2cm cube, equipped with four infrared sensors for environment sensing and communication. (b) Our Webots model implements an accurate representation of the robot and its sensors.

account the intrinsic randomness of self-assembly processes, especially at small length scales. Also, our methodological approach to the modeling and design of a swarm of self-assembling robots presents important similarities (e.g., multiple abstraction levels, assumptions about non-spatiality and sparseness at the highest level of abstraction) with the methodology we presented in a number of previous contributions (see for instance [21], [22]).

III. MATERIALS AND METHODS

This section describes the experimental setup as well as the simulation and modeling tools that we used for this case study.

A. Experimental Setup

The Alice mobile robot has a differential wheel drive that reaches speeds of 4cm/s, and four infrared sensors that allow the detection of passive obstacles at ranges up to 3cm simultaneously with 4bps bidirectional communication up to 6cm. The testing environment is a 50cm square arena.

We employ an overhead camera in conjunction with SwisTrack, an open-source object tracking tool targeted for multi-agent systems [23]. Due to our limited camera resolution, the Alice’s small size and the aggregation task itself implying continuous merging and separation of robots, markerless tracking and corresponding software analysis is not yet sufficiently reliable to use without thorough human supervision. We are currently developing enhancements that will allow SwisTrack to perform a more fully automated analysis under these conditions. The analysis of real-world experiments that follow in this paper are the result of combining SwisTrack with heavy manual verification and correction.

B. Realistic simulation

We additionally implement the above experimental setup and hardware in Webots, a realistic simulator that is able to accurately model detailed characteristics of the Alice robot, including nonlinear noisy response of the sensors and wheel slip. Although Webots provides physics-based simulations, many of these features (e.g., friction) have been disabled in the simulations presented in this paper in exchange for increased simulation speed. Webots is particularly useful




Fig. 2. Two types of discrepancies between Webots simulations and real experiments are observed: (a) in reality, a robot (encircled in red) may try to leave a chain unsuccessfully because it is blocked by its neighbors (an effect not observed in Webots), or (b) a robot may aggregate on the side of a chain due to infrared reflections that are not modeled by Webots.

because it allows us to perform fast, automatic data collection and analysis over various parameter sets. For our systematic experiments, we used a computational cluster of 50 machines, each with an Intel Pentium 4 3.00 GHz and 1GB RAM. Transitions between experiment states (i.e. changes in the total number of aggregates of each size) were stored on disk for later analysis. Our Webots simulations execute over five times faster than real time and no manual analysis is required as we have access to the internal state of the robots.

There are some discrepancies between Webots simulations and real experiments. Due to the Alice’s weak motor system, occasionally a robot may become physically trapped by others in the same aggregate (see Figure 2a), making dis-aggregation impossible. This issue is not captured in Webots as certain aspects of physical simulation have been disabled in favor of increased simulation speed.

The second observed discrepancy is that some robots may aggregate with the side of a chain due to infrared reflections that are not modeled in Webots, leading in some cases to unstructured aggregates (see Figure 2b), a phenomenon also observed in previous aggregation experiments [11]. The probability of occurrence of this type of aggregation is a function of the chain size (the more robots in a chain, the more in-chain binding sites) as well as the shape and sensor arrangement of the robots. As a result, this effect is very difficult to capture accurately in a realistic simulation like Webots, and manually collecting enough experimental data in order to obtain a statistically relevant measure of its probability of occurrence is impossible at this time. As such, this discrepancy is currently ignored by our models.

C. Baseline Aggregation-Only Deterministic Controller

In all of our controllers, the aggregation behavior of the robots is the same: they perform a random walk and, upon collision, may decide to aggregate as a function of their relative alignment (Figure 3), which can be roughly estimated using short-range proximity sensors. In this baseline controller, we assume that the robots, once aggregated, never disaggregate. As we will see later, this behavior invariably

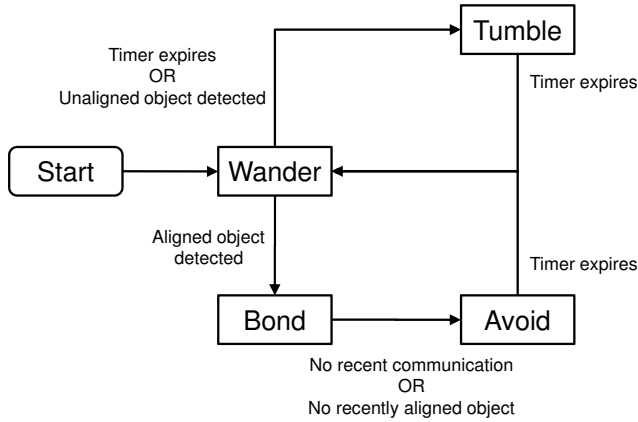


Fig. 3. Basic algorithm for aggregation and random walk. The robots begin by moving in a straight path. After some period of time, or if an unaligned object is encountered, it begins a random tumble, afterwards returning to straight line travel. In this way we achieve a well mixed system in which robots “bounce” away from walls and unaligned Alices. Should a well-aligned object be encountered, the robot stops in place and attempts to communicate. If communication is not returned, or the object leaves, a brief period of avoidance is entered before returning to the initial state.

leads to an exponential-like distribution of chain sizes, which we call the *basic distribution*.

D. Deterministic and Probabilistic Disaggregation Controllers

We implement two controllers based on the above algorithm, which differ only in the way they handle disaggregation.

Our first extension implements a deterministic controller, which will choose not to aggregate with chains already equal to the target length. Each robot communicates on one sensor an incremented observation of the opposite side (see Figure 4). Thus, outer nodes will send a value of zero towards the center of a chain. This number is incremented as it passes through each robot in the chain, eventually communicating its total size to the outer nodes. Other Alices that attempt to bond at the ends of a full chain will receive this value, causing them to fail to aggregate.

Implementing a controller based on local communication was not easy on robots of this size, as depending on lighting conditions and robot density, the Alices may experience extreme packet corruption and loss. In our algorithm, this will occasionally lead to a significant increase of the amount of time a chain’s size takes to propagate to its outer nodes, allowing undesirable aggregations to occur in the interim. While this is a rare occurrence with the Alice robots (and as such not captured in our models), we believe that such issues with local communication will become more severe as our target systems decrease in scale.

In contrast, we implement an alternative controller that requires no explicit communication between robots. Instead of limiting the maximum size of a chain directly, in this method Alices leave their current bond with probability $p_{leave,1}$ or $p_{leave,2}$ each second, depending on whether it





Fig. 4. A typical example of local communication in a system using the deterministic controller and a target chain size of four: (i) one robot joins a chain while the robot at the other end sends a value of zero into a chain. (ii) The adjacent robot reads an integer on one sensor, increments it, and (iii) communicates it on the opposite side. As time progresses, each robot learns its position within the chain and transmits to its next neighbor. (iv) Finally, the robot that has just aggregated receives a message equal to the chain size and (v) does not aggregate because this value is too large.

has one or two neighbors, respectively. In other words, Alices in a pair or at the end of a chain will leave with probability $p_{leave,1}$, while robots in the middle of a chain will leave with probability $p_{leave,2}$. Part of this paper explores how one can control the length of chains when the only input parameters are these leaving probabilities.

E. Macroscopic Modeling

Hereafter, we describe a macroscopic model of the two above controllers where interactions between robots are modeled as a Chemical Reaction Network (CRN). A set of reactions can be represented as a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. The set of vertices \mathcal{V} represents the *complexes*, i.e. a pair of aggregating robots or one robot aggregating with a chain. The set of directed edges \mathcal{E} represents the *reactions* between complexes, i.e. the aggregation and disaggregation of robots. Each reaction can be denoted by an ordered pair $(i, j) \in \mathcal{V} \times \mathcal{V}$, meaning that one complex i transitions to complex j . Furthermore, each reaction is associated with a rate constant.




We consider the quantity of single robots X_s and chains X_i , where i is the number of robots in the chain. As chains remain unmoving in the environment, only a few complexes are possible in our system: $X_s + X_s$, which denotes the collision of two single robots, and $X_s + X_i$, which denotes the collision between a single robot and a chain of size i . As a result, formulating the set of *forward reactions*, i.e. those corresponding to an aggregation event, is straightforward:





where N_0 is the total number of robots, and e_s and e_i are the rates at which a robot aggregates with another single robot or a chain of size i , respectively.

On the other hand, the *backward reactions*, i.e. those corresponding to a disaggregation event, are slightly more difficult to define. Indeed, for a chain of size i , any robot

may potentially leave, thus leading to i different reactions:



\vdots



Now, it is clear that Equation 3 and Equation 7 are the same, as are Equation 4 and Equation 6. In fact, each pair of symmetric equations can be merged into a single equation with a doubled rate constant. As a result, there are $i/2$ and $(i + 1)/2$ backward reactions for a chain of size i with i even and odd, respectively (note that the reaction corresponding to the disaggregation of a robot in the middle of an odd-sized chain has no symmetrical counterpart). Therefore, the number of possible reactions in a system composed of N_0 robots is $O(N_0^2)$. For small N_0 , it is possible to model the entire CRN; however, when N_0 is large, the number of reactions explodes and the system is no longer tractable. As a result, we assume that chains larger than a given upper bound N_{max} cannot be formed, thus limiting the number of reactions. This assumption is perfectly valid for sufficiently large N_{max} as the probability that very large chains are formed is negligibly small. Note that the bound N_{max} can be trivially used for modeling the aforementioned baseline controller by setting the rate constant of the backward reactions to zero.

We take $k_{i,j}^-$ to be the rate of disaggregation of a chain of size i due to the leaving of its j -th robot. As stated before, we merge symmetric reactions into a single reaction with $k'_{i,j} = 2 \cdot k_{i,j}^-$ with $j = 1, \dots, i/2$ if i is even and $j = 1, \dots, (i - 1)/2$ if i is odd. Recall that the disaggregation of a robot in the middle of an odd-sized chain has no corresponding symmetric reaction, and as such its rate constant remains unchanged.

We calculate the rate constants themselves using geometric approximations. If we assume that our system is well-mixed¹ we can capture the aggregation process with a nonspatial model. Therefore, we assume that the probability a robot is at a given position is independent of time and uniformly distributed over the arena space. Thus, the probability of two specific robots encountering each other during a discrete time interval T is given by

$$p_c \sim \frac{\hat{v} T w_d}{A_{tot}} \quad (8)$$

where \hat{v} is the average velocity of a robot, $w_d = 2r$ is the diameter of a robot, and A_{tot} is the total area of the arena [21].

¹As robots form chains, the chains themselves may significantly constrain the accessibility of different parts of the arena. We ignore this effect in our macroscopic models, which will limit their real-world correlation if the arena used is not sufficiently large.




Fig. 5. In our models, we assume that two colliding robots can assemble only if their relative heading θ_h is smaller than $\alpha_a/2$. The blue circular sector represents both the detection and the communication area of the robots.

A collision does not necessarily lead to aggregation, due to the directionality characterizing our building blocks: robots must be aligned in order to aggregate. More formally, this means the absolute value of the heading θ_h must be smaller than $\alpha_a/2$, where α_a is the central angle of the detection and communication sector (see Figure 5). Because of the non-holonomic nature of the robots, we assume that there is always at least one robot, which we denote **B**, that is aligned upon collision (**B** is basically the robot that runs into the other). Furthermore, we assume that the absolute value of the heading of **B** with respect to **A** is uniformly distributed in $[0, \pi]^2$. As a result, the probability p_a that, upon collision, two robots are properly aligned is $p_a \sim \alpha_a/(2 \cdot \pi)$. Since each robot and each chain has two valid binding sites, the rate constants e_s and e_i can be written

$$e_s = e_i = p_c \cdot 2 p_a \cong \frac{\hat{v} T w_d}{A_{tot}} \cdot \frac{\alpha_a}{\pi} \quad (9)$$

In our first approximation, the rate constants of the backward reactions do not depend on the geometry of the robots, but on the leaving probabilities encoded in the controller. We consider two cases: (1) a robot has a single neighbor and his leaving probability is $p_{leave,1}$, or (2) a robot has two neighbors and his leaving probability is $p_{leave,2}$. As we observed in real experiments, some robots may be unable to leave because they are physically trapped between their neighbors. This effect can be captured by decreasing the leaving probability by a certain factor $1 - p_s$, where p_s corresponds to the probability that a robot remains stuck. As a result, the rate constants for the backward reactions can be written

$$k_{i,1} = p_{leave,1} \quad (10)$$

$$k_{i,j} = p_{leave,2} \cdot (1 - p_s) \quad \text{for } j \neq 1 \quad (11)$$

Since there is no simple geometrical approximation for the probability p_s , we need to measure it using real experiments. In our case, we assume that p_s is negligibly small.

²Note that because single robots are always moving forward, it is more likely for **B** to encounter the front of **A** than its rear. We do not capture this in our models, although including this distribution may be an interesting enhancement in future work.

This CRN can be easily abstracted to an ODE model using the relation $\dot{\mathbf{y}} = \mathbf{S} \cdot \mathbf{p}(\mathbf{y})$ where $\mathbf{S} = (s_{ij})$ is the stoichiometry matrix, with the stoichiometric coefficient s_{ij} of the j -th species in the i -th reaction, and $\mathbf{p}(\mathbf{y})$ is the propensity vector, which depends on the reaction rates and the population of reactants for each reaction. This system can be solved using conventional integration techniques. We know that in the thermodynamic limit, i.e. when the number of robots tends towards infinity, the ODE approximation converges to an exact solution. However, since our system contains only $N_0 = 19$ robots, the ODE representation introduces an approximation error which is in our case negligible (see Section IV-A).

IV. RESULTS

We implemented the two types of controllers on a real swarm of 19 Alice robots and verified the qualitative relevance of our models and simulations in five real experiments of 20 minutes. Using both the deterministic and the probabilistic controllers, we observed the formation of chains of variable length in the system. Based on these systematic experiments, we could validate both our Webots and CRN models (Figure 8). We also carried out non-systematic real experiments with different initial conditions and duration that are not reported in detail hereafter. A 16 minute excerpt from one such experiment is provided in Figure 6 (this particular experiment can be seen in full in the included video).

A. Macroscopic models validation

As mentioned previously, the ODE approximation is guaranteed to be valid in the thermodynamic limit, but not necessarily with $N_0 = 19$. Therefore, in order to validate our ODE model, we compare it with stochastic simulations of the CRN, such as the Gillespie method (exact simulation), or its optimized counterparts (approximate simulation), provided by the StochKit toolbox [24]. In fact, the ODE model exhibit an excellent accuracy for $N_0 = 19$, well within the standard deviation of stochastic simulations (Figure 7).

In order to validate our macroscopic models, we use realistic Webots simulation as a baseline. Figure 9 compares the prediction of stochastic simulations and Webots simulations in the case of the probabilistic controller with $p_{leave,1} = 10^{-4}$ and $p_{leave,2} = 10^{-9}$. We see an excellent fit for chains of size four, however smaller chains tend to grow faster in stochastic simulations than in Webots due to abstractions discussed in Section III-E.

B. Deterministic controller simulations

We performed a series of Webots simulations (averaged over 100 runs) and stochastic simulations (averaged over 500 runs) for the deterministic controller with varying maximal chain sizes from two to eight with $N_0 = 19$ robots (Figure 10). These results demonstrate that the controller successfully limits the size of the chain, but, as the target chain size increases, the distribution of chain size approaches the one yielded by the baseline aggregation-only controller (see Section III-C), which is depicted in Figure 11 for $N_0 = 19$ robots .




Fig. 6. 16 minute excerpt from a real experiment with $N_0 = 15$ robots and $p_{leave,1} = p_{leave,2} = 10^{-3}$. (a) At $t = 0$, the system is in an arbitrary, well mixed state. (b) After two minutes, some aggregates have already begun to form. (c), (d), (e) The system continues to evolve from four minutes to 12 minutes; images captured at uniform intervals. (f) Finally, at 16 minutes, most Alices have bonded with others to form an aggregate (only a single robot remains disaggregated).

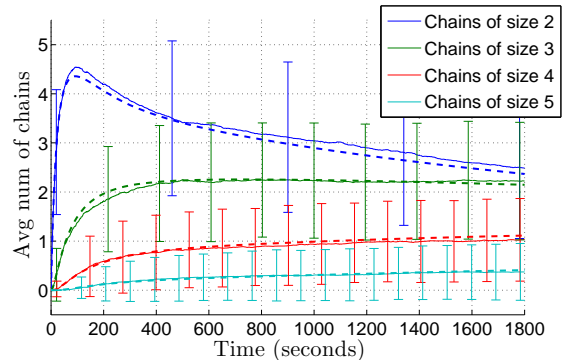


Fig. 7. Comparison of the prediction of stochastic simulations (continuous lines) and the ODE model (dashed lines) for the time evolution of the system with $N_0 = 19$, $p_{leave,1} = 10^{-4}$ and $p_{leave,2} = 10^{-9}$. These results show that, for our case study, the ODE approximation remains valid even for small N_0 .




Fig. 8. Probabilistic controller validation (black squares are experimental data averaged over 5 runs) of CRN models and Webots simulations using the average number of chains of size three over the course of the experiments. $N_0 = 19$, $p_{leave,1} = p_{leave,2} = 10^{-3}$. Large variance is observed due to all techniques being highly stochastic in conjunction with sensor and actuator noise in both Webots and the real robots.




Fig. 9. Comparison of the probabilistic controller's time evolution of the average number of chains of size two, three and four predicted by stochastic simulations (dashed lines) and Webots simulations (continuous lines). $N_0 = 19$, $p_{leave,1} = 10^{-4}$, $p_{leave,2} = 10^{-9}$.

C. Probabilistic controller simulations

We performed a series of Webots simulations over the entire parameter space $(p_{leave,1}, p_{leave,2}) \in [10^{-9}, 10^{-1}]^2$ (9×9 logarithmic discretization, averaged over 100 runs for each parameter set) with a simulated swarm of 19 Alice robots. Trials lasted 30 minutes of simulated time. In a parallel effort, we simulated the CRN presented in Section III-E using StochKit over the same time span and parameter space, but with a finer discretization of 16×32 averaged over 500 runs.

The results of these searches, depicted in Figure 12, confirm the existence of an optimal region for $p_{leave,1} \in [10^{-4}, 10^{-3}]$ and $p_{leave,2} < 10^{-4}$ where the ratio between the number of pairs and the number of longer chains decreases because exploration (disaggregation) and exploitation (aggregation) are well balanced, thus leading to a chain size distribution different than the basic distribution.

The long term evolution of the distribution of chain size for leaving probabilities that lie within the optimal region ($p_{leave,1} = 10^{-4}$ and $p_{leave,2} = 10^{-9}$) demonstrates clearly that one can achieve a basic control over the distribution of chain size by setting appropriate leaving probabilities only. More precisely, it is possible to skew the steady state distribution towards longer chains. Higher, non-optimal leaving probabilities enable faster convergence, but they do not favor the formation of long chains.




Fig. 10. Purely deterministic control does not yield desirable results when attempting to form long chains. As we increase the maximum chain size from two to eight with $N_0 = 19$ robots, we show in both (a) Webots (averaged over 100 runs) and (b) stochastic simulations of the CRN (averaged over 500 runs) that a lack of exploration causes the system to quickly converge to the basic distribution, in which short chains are largely favored.




Fig. 11. Comparison of the basic distribution for $N_0 = 19$ robots using Webots (dark grey) and stochastic simulations (light grey). It is clear that setting leaving probabilities to zero does not favor the formation of long chains, thus some amount of exploration is required.

D. Discussion

Our results have shown that one cannot favor the formation of chains of a specific size only by limiting the maximum chain size (deterministic controller) or by setting specific leaving probabilities (probabilistic controller). Indeed, the landscapes depicted in Figure 10 and Figure 12 do not exhibit a systematic variation in the ratio of different chain size populations that would enable a precise selection of a particular chain length. On the one hand, the deterministic controller, by using communication among the robots, can achieve a proper non-linear feedback that depends on the size of the chain, thus preventing the formation of chains longer than a certain target size, whereas the probabilistic controller is purely linear, i.e. the robots have the exact same behavior regardless of the size of the chain. On the other hand, the deterministic controller is unable to achieve long chains because of its intrinsic lack of exploration whereas the probabilistic controller enables, with proper leaving probabilities, a balance of exploration and exploitation. One potential solution to this problem could be to combine the deterministic and the probabilistic controllers into a single, hybrid controller: indeed, by optimizing the leaving probabilities

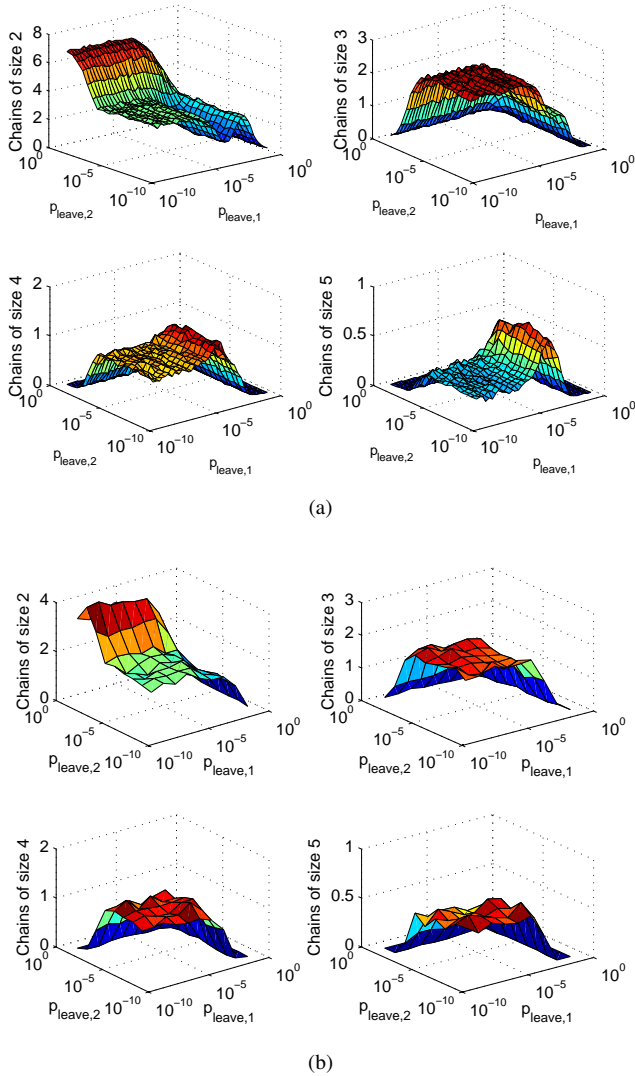


Fig. 12. Systematic exploration of the probability space spanned by $p_{leave,1}$ and $p_{leave,2}$ (logarithmic scale) for $t = 30\text{min}$. (a) The existence of an optimal region for $p_{leave,1} \in [10^{-4}, 10^{-3}]$ and $p_{leave,2} < 10^{-4}$ where exploration (disaggregation) and exploitation (aggregation) are well balanced is confirmed by stochastic simulations. (b) The lower resolution and the larger noise of Webots simulations make this optimal region fade out, but the general landscape is qualitatively the same.

for the formation of arbitrarily long chains while limiting explicitly their size to a given upper bound N_{max} , one would most likely end up favoring the formation of chains of size N_{max} . However, this hybrid approach would be less scalable than the purely probabilistic one presented in Section III-D or a timeout-based approach [21].

An interesting insight provided by our results is that the self-assembly process takes a great amount of time to stabilize, in particular when leaving probabilities are low, as depicted by Figure 13. These results demonstrate clearly that, for small leaving probabilities, the time scale of the self-assembly process is much larger than that of the experiments presented in this paper. As a result, the influence of the basic distribution is very important in the beginning, as it is essentially the underlying distribution of any short term experiments, and it decreases with time. This apparent flaw

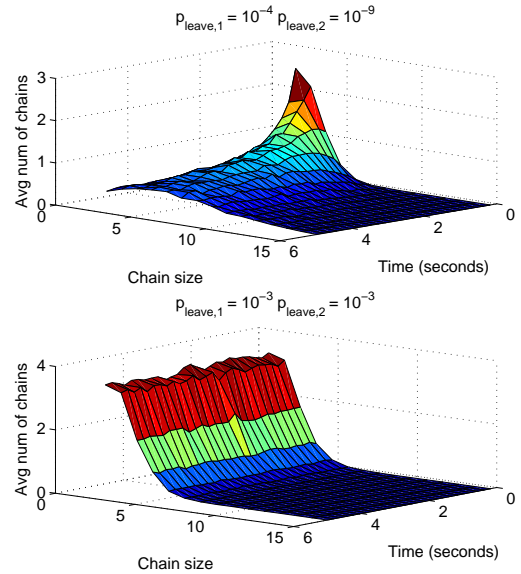


Fig. 13. Long term prediction of the chain size distribution up to $t_f = 5 \cdot 10^4 s$ (about 14 hours) for two different parameter sets (stochastic simulations). For small probabilities ($p_{leave,1} = 10^{-4}$ and $p_{leave,2} = 10^{-9}$), it is clear that the convergence is very slow and the time required to reach the steady state is at least one order of magnitude higher than the duration of our experiments. For larger probabilities ($p_{leave,1} = 10^{-3}$ and $p_{leave,2} = 10^{-3}$), the convergence is much faster and on the same time scale as the duration of our experiments (about 30 minutes).

of the approach is actually a great opportunity: by choosing a proper duration of the process, one can actually control how close to the basic distribution the system will be. Of course, this approach is possible only if the leaving probabilities depend on the time (in a very basic fashion, since they will simply be set to zero after a time t_f)³. Therefore, the control parameters of the system are not only $p_{leave,1}$ and $p_{leave,2}$, but also the final time t_f , yielding a 3-dimensional parameter space. It is clear that a fine-grained systematic search over this entire space is prohibitively expensive from a computational standpoint. It is intractable with Webots simulations, and perhaps even with macroscopic models, if solved exactly with their stochastic interpretation, a high-performance computational cluster would be required, both for simulation and data analysis. Instead, we plan to use standard nonlinear optimization techniques based on either the deterministic formulation of the system (ODE) in its desired steady state or a reduced linearized formulation of the CRN [25]. We will explore these solutions in future work.

V. CONCLUSION AND FUTURE WORK

In this paper, we proposed two approaches to the control of a swarm of self-assembling miniaturized robots: (1) a deterministic controller, non-linear in the robot cluster size, where robots communicate with each other in order to determine the size of the chain and adapt their behavior accordingly,

³Note that this requirement does not jeopardize the scalability of our approach since one could imagine to use an external observer to broadcast a predetermined message or to modify an environmental parameter of the system (e.g., illumination, temperature, pH) at t_f in order to signal the termination of the self-assembly process.

and (2) a probabilistic controller, linear in the robot cluster size, where the distribution of chain size is controlled by the leaving probabilities of the robots. We demonstrated the feasibility of both approaches by implementing them on a real swarm of 19 Alice mobile robots. Using a realistic robotics simulator in parallel with macroscopic models based on the Chemical Reaction Network (CRN) framework, we performed systematic searches of the parameter space that demonstrated (1) the inability of the deterministic controller to achieve long chains because of its lack of exploration and (2) the existence of optimal regions within the parameter space of the probabilistic controller where exploration and exploitation are well balanced, thus favoring the formation of larger chains. However, the linearity in robot cluster size of the probabilistic controller prevents any positive feedback that would favor specifically a given target chain size. A potential solution to this problem is to develop a hybrid controller with optimized leaving probabilities for the formation of arbitrarily long chains while explicitly limiting their size to a given target size N_{max} .

As mentioned before, one of the long term objectives of our work is the development of scalable control methodologies for large swarm of miniaturized robots. In this paper, we achieve a first step towards this goal by proposing a probabilistic controller that has the potential to scale down in terms of size of the robot, since it is fully reactive and memoryless, while being intrinsically scalable in terms of size of the swarm, since it is completely decentralized. We plan to pursue this research by applying the methodological framework presented in this paper to the fluidic self-assembly of Micro-Electro-Mechanical Devices (MEMS) building blocks, which we believe are the precursor to actual micro-robots endowed with basic sensory, actuating, and computational capabilities. However, two major obstacles lie on the route to this objective: (1) we must be able to better capture spatiality, embodiment and suboptimal mixing of the targeted systems at the macroscopic level, and (2) we need to update our experimental setup in order to enable automated analysis of our experiments, at all scales. Solving these problems are major research efforts that we are currently pursuing.

VI. ACKNOWLEDGMENTS

The authors would like to thank Loïc Matthey for helpful discussion and useful scripts for analyzing stochastic simulations, and our anonymous reviewers for their thorough comments.

REFERENCES

[1] A. Zangwill, "Statistical physics: Advances in aggregation." *Nature*, vol. 411, no. 6838, pp. 651–652, 2001.
 [2] T. Clark, J. Tien, D. Duffy, K. Paul, and G. Whitesides, "Self-assembly of 10-m-sized objects into ordered three-dimensional arrays," *J. of the American Chemical Society*, vol. 123, pp. 7677–7682, 2001.
 [3] J. Parrish and W. Hamner, *Animal Groups in Three Dimensions: How*

Species Aggregate. Cambridge University Press, 1997.
 [4] W. Agassounon, A. Martinoli, and K. Easton, "Macroscopic modeling of aggregation experiments using embodied agents in teams of constant and time-varying sizes," *Autonomous Robots*, vol. 17, no. 2-3, pp. 163–192, 2004.
 [5] R. Gross, M. Bonani, F. Mondada, and M. Dorigo, "Autonomous self-assembly in swarm-bots," *IEEE Trans. on Robotics*, vol. 22, no. 6, pp. 1115–1130, 2006.
 [6] G. Whitesides and B. Grzybowski, "Self-assembly at all scales," *Science*, vol. 295, no. 5564, pp. 2418–2421, 2002.
 [7] G. Caprari and R. Siegwart, "Mobile micro-robots ready to use: Alice," *Intelligent Robots and Systems, 2005. (IROS 2005). 2005 IEEE/RSJ International Conference on*, pp. 3295–3300, 2005.
 [8] D. Wilkinson, *Stochastic Modelling for Systems Biology*, ser. Mathematical and Computational Biology Series. CRC Press, Jan 2006.
 [9] D. Gillespie, "Exact stochastic simulation of coupled chemical reactions," *The J. of Physical Chemistry*, vol. 81, no. 25, pp. 2340–2361, Jan 1977.
 [10] —, "Stochastic simulation of chemical kinetics," *Annual Review of Physical Chemistry*, vol. 58, pp. 35–55, 2007.
 [11] A. Martinoli, "Swarm intelligence in autonomous collective robotics: From tools to the analysis and synthesis of distributed control strategies," Unpublished doctoral manuscript n. 2069, EPFL, Lausanne, October 1999.
 [12] S. Garnier, C. Jost, J. Gautrais, M. Asadpour, G. Caprari, R. Jeanson, A. Grimal, and G. Theraulaz, "The embodiment of cockroach aggregation behavior in a group of micro-robots," *Artificial Life J.*, vol. 14, no. 4, pp. 387–408, Jan 2008.
 [13] J. Halloy, G. Sempo, G. Caprari, C. Rivault, M. Asadpour, F. Tache, I. Said, V. Durier, S. Canonge, J. Ame, C. Detrain, N. Correll, A. Martinoli, F. Mondada, R. Siegwart, and J. Deneubourg, "Social integration of robots into groups of cockroaches to control self-organized choices," *Science*, vol. 318, pp. 1155–1158, 2007.
 [14] E. Klavins, "Programmable self-assembly," *IEEE Control Systems Mag.*, vol. 27, pp. 43–56, 2007.
 [15] L. Matthey, S. Berman, and V. Kumar, "Stochastic strategies for a swarm robotic assembly system," in *IEEE 2009 International Conference on Robotics and Automation (ICRA 2009)*. IEEE, 2009, pp. 1953–1958.
 [16] S. Berman, A. Halasz, V. Kumar, and S. Pratt, "Bio-inspired group behaviors for the deployment of a swarm of robots to multiple destinations," *Robotics and Automation, 2007 IEEE International Conference on*, pp. 2318–2323, 2007.
 [17] A. Julius, A. Halasz, M. Sakar, H. Rubin, V. Kumar, and G. Pappas, "Stochastic modeling and control of biological systems: The lactose regulation system of *escherichia coli*," *IEEE Trans. on Automatic Control*, vol. 53, pp. 51–65, Jan 2008.
 [18] H. Tanner, A. Jadbabaie, and G. Pappas, "Flocking in teams of nonholonomic agents," *Lecture Notes in Control and Information Sciences*, vol. 309, pp. 229–239, 2005.
 [19] A. Jadbabaie, J. Lin, and A. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Trans. on Automatic Control*, vol. 48, no. 6, pp. 988–1001, 2003.
 [20] G. Pereira, V. Kumar, and M. Campos, "Closed loop motion planning of cooperating mobile robots using graph connectivity," *Robotics and Autonomous Systems*, vol. 56, no. 4, pp. 373–384, 2008.
 [21] A. Martinoli, K. Easton, and W. Agassounon, "Modeling swarm robotic systems: A case study in collaborative distributed manipulation," *Int. J. of Robotics Research*, vol. 23, no. 4-5, pp. 415–436, 2004.
 [22] G. Mermoud, J. Brugger, and A. Martinoli, "Towards multi-level modeling of self-assembling intelligent micro-systems," in *AAMAS '09: Proc. of the 8th Int. Conf. on Autonomous Agents and Multiagent Systems*, Richland, SC, May 2009, pp. 89–96.
 [23] T. Lochmatter, P. Roduit, C. Cianci, N. Correll, J. Jacot, and A. Martinoli, "Swistrack - a flexible open source tracking software for multi-agent systems," in *IEEE/RSJ 2008 International Conference on Intelligent Robots and Systems (IROS 2008)*. IEEE, 2008, pp. 4004–4010.
 [24] H. Li, Y. Cao, L. Petzold, and D. Gillespie, "Algorithms and software for stochastic simulation of biochemical reacting systems," *Biotechnology Progress*, vol. 24, no. 1, pp. 56–61, Jan 2008.
 [25] L. Matthey, "Hybrid reactions modeling for top-down design hybrid reactions modeling for top-down design framework," Master's thesis, EPFL, KUMAR-SWIS-MP12, 2008.