


## Model-Based Control of Programmable Self-Assembly of Lily Robots

Maximilian Mordig

Professor : Alcherio Martinoli

Assistant(s) : Bahar Haghighat

Stochastic self-assembly is defined as the reversible phenomenon of an ordered spatial structure emerging from the aggregate behavior of simpler entities through local interactions and inherent randomness, with no external direction. Self-assembly can be observed at all scales, from molecules to weather phenomena. Understanding how to steer such systems from an initial configuration to some target configuration is crucial to leverage self-assembly to build structures. The project considers Lily robots to build chains of Lilies starting from single Lilies. These Lily robots have local rulesets that allow to accept or reject to connect to other Lily robots. A chain of length  $n$  corresponds to an  $n$ -mer. The goal of this project is to improve the prediction of the copy numbers of each  $n$ -mer over time. The data used to model the system was collected with Webots as shown in the figure.

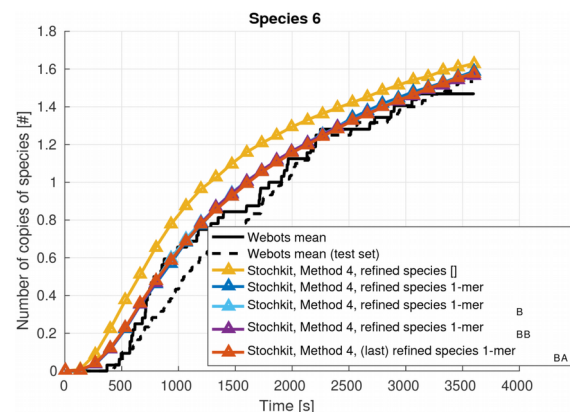


The system is modeled as a Chemical Reaction Network (CRN) with species, reactions and reaction rates. The species are the  $n$ -mers. The macrostate consists in the copy number of each  $n$ -mer. The predictions are shown in the figure on the right for species 6 (6-mer). The predictions by the initial CRN are good, yet they can be improved by “refining” the initial CRN as shown in the same figure. The refinement of a CRN replaces a species by two species, adds corresponding reactions and uses an expectation-maximization algorithm to infer the new species populations from the observed macrostates.

To do so, the code was rewritten from a previous project and entirely restructured. The code is very modular, e.g. the refinement procedure can be adapted to add other reactions or refine a species into

more than two species. The project had a large software engineering aspect.

The next contribution was to propose a rule on how to select the next species to refine as well as a stopping criterion. It was observed that the first refinement had the greatest impact and it was most important to refine the first species (1-mer), probably because it occurs most often. Also, it was observed that the refinement procedure is non-unique meaning that different CRN can give rise to very similar species trajectory predictions.



During the course of the project, it was detected that the data collected by the Webots simulations is inconsistent and this same inconsistent data was also used in previous projects. Attempts to correct this were unsuccessful. This questions some of the results of the project. Consequently, fake data was generated to check the correct working of the refinement code successfully.

Finally, the last part of the project considered a system with two different modes, each mode governed by a different ruleset (but the same ruleset for all Lily robots). The goal was to choose an optimal mode at each time step to steer the system as quickly as possible to the target configuration. For this, the previous part was used to model the system as two (initial) CRNs, one CRN for each mode. Then, the problem was written as a Bellman equation corresponding to a Markov decision process and solved by Asynchronous Value Iteration. It was shown that in the presence of two modes, the overall assembly time from the starting to the target population can be reduced by 10% and it is hence beneficial to use multiple modes.