Body Building: Hatching Robot Organisms

Berend Weel Evert Haasdijk A.E. Eiben

***Abstract*—We propose a new scenario for the evolution of robot morphologies based on an egg metaphor. A swarm of robots is released in a large arena in which they form organisms through a process of morphogenesis. These organisms can reproduce by fertilising eggs in their vicinity, these fertilised eggs in turn build new organisms by recruiting free modules as a 'seed'. We investigate the influence of three parameters of this evolutionary system: the time eggs wait to be fertilised, the maximum time a seed can recruit modules to form an organism and how long an organism lives. Specifically we investigate the influence of these parameters on the size and stability of the population of eggs, seeds and organisms. It is shown that the influence of the lifetime of an organism is the largest, and leads to many organisms, it should be set much higher than the other two. Furthermore setting the time an egg can be fertilised and the maximum time a seed is allowed to build an organism to the same value results in the most stable system regarding number of eggs and seeds. Finally setting the maximum seed time higher leads to to a slightly smaller population of bigger organisms.**

***Index Terms*—Embodied Evolution, Robotic Organisms, Evolutionary Robotics, Robotics, Robot Morphology**

I. Introduction

The majority of research in the field of evolutionary robotics focussed on the development of robot controllers in an offline manner, i.e. with a centrally controlled evolutionary algorithm that runs on an external computer. The result of this evolutionary algorithm is transferred to real robots after the run has finished. In some cases the controllers created by this central evolutionary algorithm are tried in a simulator, e.g. fully offline, while in other cases the trials are in real robots called 'Embodied Trials'. In the paper by Watson et al. [18], a new paradigm called Embodied Evolution is proposed: one where the evolutionary algorithm runs on the robots themselves as they go about their tasks.

We can take embodiment one step further as stated by Eiben et al. [6]: the evolution of the bodies themselves. This form of Embodied Evolution could be the next step in evolutionary computing, going from an evolutionary process and result in digital space, to an evolutionary process and result in physical space. To enable the evolution of robot bodies, or morphology, we need to have a medium for evolution. There are, in our view, two options: 3D printing and Modular Robotics, and in this paper we focus on the latter option.

Consider a scenario where a number of robotic organisms populate an arena and perform some task. Each organism consists of several robot modules, and each has its own distinctive shape. These organisms move and perform the task, although some are more adapted to the environment and task than others.

Each organism carries a genome that encodes its shape. To reproduce, organisms need to find an 'egg': an individual stationary robot module that they can 'fertilise' with their genome. An egg that has been fertilised by a number of organisms selects parents from the received genomes for recombination and/or mutation. The resulting new genotype is the basis for a new organism that will be born through a process of morphogenesis that the egg initiates. The egg becomes a seed, recruiting available modules in order to build the required shape; when this process is complete, a new organism is born.

This new organisms body is slightly different from that of its parent(s), so it has to learn how to use this body effectively through lifetime learning. It takes a while until the controller has adapted to control its body, but after some time the controller is able to move and perform the task with the new bod and try to propagate its genes.

After a certain amount of time has passed an existing organism comes to the end of its life. The organism disassembles and individual modules detach. Most of the modules drive away in search of a forming organism that they will become part of, one of them stays put and becomes an egg to be fertilised.

The scenario envisioned above has three major parameters: how long eggs listen for genomes (*egg lifetime*), how long a fertilised egg (a seed) is allowed to build the organism its genome encodes (*seed lifetime*), and finally how long a fully grown organism is lives before it dies (*organism lifetime*).

To understand the 'inner logic' of this system we conduct experiments to see how these parameters interact, in particular regarding their influence on:

- The size of the organism population;
- The stability of the organism population;
- The average size of the organisms.

Besides gaining insight, the goal of this investigation is similar to that of [5]; we are looking for a set of parameters that leads to a large, stable population of adult, i.e. completed, organisms, as well as reasonable and stable amount of eggs that can be fertilised. For now we investigate these questions in simulation as a proof of concept and to quickly get insights into the dynamics of the system.

A. Related Work

Evolutionary Robotics has been widely studied since the early 1990s [12]. One avenue of research uses evolution in large numbers of interacting robots in swarms [4] or modular robots (for examples see [20]). In these cases, evolution is typically used to achieve some fixed user-defined objective, such as locomotion or explicit coordination.

The scenario we envision can be implemented with and without a user specified selection criterion. Open-ended or objective-free evolution (evolution without a user specified selection criterion) has been studied in Artificial Life since Rasmussen's [15] and Ray's [16] work. The primary goal of this research is to investigate evolutionary dynamics in the absence of tasks. Instead of explicit task performance it is the ability to spread genomes through the population that drives evolution, leading to an implicit, or environmental, selection pressure. This open-ended approach has gained interest from the evolutionary robotics community, for instance in Bianco and Nolfi's experiments with self-assembling organisms [1] and more recently in the MEDEA algorithm [2].

When evolving the morphology of modular robots these modules need to self-assemble. Self-assembly has been demonstrated in numerous cases: [7], [10], [13], [19], [20]. The process of morphogenesis has been identified as a key feature for the SYMBRION project [9], and has been investigated in several instances [3], [8], [11], [14], [17].

II. System Description

We have implemented the scenario detailed in the introduction in a simple 2D simulator: RoboRobo[1] using simulated e-puck robots. Of course, in a 2D environment the movement of a robot organism is much simpler than in a 3D setting, however for the topics investigated in this paper it suffices. The robots can steer by setting its speed and desired rate of turn and each robot has 8 sensors to detect obstacles (static obstacles as well as other robots).

The scenario we created takes place in a large empty arena which is 64 robots wide on each side. We divide a swarm of 100 robots into two groups: eggs and free modules. Each egg can potentially grow into an organism using its genome, while the free modules always play an assisting role and do not have a genome for an organism shape. In our experiments we used 20 eggs and 80 free modules, meaning that we can have at most 20 organisms at the same time. We also set a maximum organism size of 10. Obviously we cannot have 20 organisms of size 10 at the same time, however this is of minor issue as the goal is to have both living organisms of different sizes and eggs listening for genomes at the same time. At the start of our experiment eggs are initialised with a random genome and half of them are set to be seeds that immediately start to create the organisms encoded by their genomes. The remaining half of the eggs remain as-is and listen for genomes. To diminish the periodicity of the system this first generation of eggs and seeds is initialised with a random age between 0 and the maximum lifetime.

All robot modules are programmed with a simple obstacle-avoidance behaviour coupled with a random trajectory when no obstacles are near. An organism moves in the direction of the combined vectors of its constituent modules. This behaviour is reasonable in this scenario as the robots explore the entire arena and do not get stuck behind each other. It also prevents the problems of having to learn a behaviour, which is beyond the scope of this paper.

When a free module comes within a certain distance of a seed (slightly over 3 robot diameters in this case), the module is automatically drafted into the seed's organism. It is teleported to the correct position as indicated by the seeds' genome,

[1]https://code.google.com/p/roborobo/

provided this does not cause a collision. This teleportation is a shortcut for a more sophisticated docking procedure that should be implemented on real robots, but for these simulations we assume that when robots come this close to each other docking is always successful. Normally e-pucks cannot connect to each other, however we have modified our simulation to allow them to make rigid connections on four sides as an emulation of the robots used in the Symbrion project[2]. When a seed has completed its organism, the organism starts to move and spread its genome to eggs it comes close to. If a seed is unsuccessful in building its shape within the allotted time it reverts to the egg role and listens for genomes again for the prescribed time.

Eggs are automatically fertilised by organisms that come within range (again, slightly over 3 robot diameters); this introduces an environmental pressure towards organisms that move about effectively. In this scenario the pressure towards any particular kind of morphology is very weak. There is a slight favour towards smaller organisms, as it is easier to coordinate locomotion between few modules than with many. This setup has been specifically created to isolate the influence of the parameters under investigation.

Offspring: At the end of its listening period an egg produces offspring based on the genomes that it received as follows. If an egg has not received any genomes from passing organisms it uses its own genome to create offspring by mutation. If an egg has received a single genome it uses recombination between its own genome and the received genome and then mutation to create an offspring. If the egg has received two or more genomes it randomly selects two of the received genomes and then uses recombination and mutation to create the offspring. The egg becomes a seed and starts building the new organism.

Experiments: To investigate the influence of egg, seed and organism lifetime we ran a series of experiments using three values for each parameter: 1000, 5000 and $10{,}000$ time steps of our simulator, resulting in 27 different parameter sets. Each run lasted a total of $100{,}000$ time steps was repeated 40 times.

An overview of the experiment parameters can be found in Table I.

[2]http://www.symbrion.eu

Table I: Overview of experimental parameters

Parameter	Value
Egg Lifetime	1000 - 10,000
Seed Lifetime	1000 - 10,000
Organism Lifetime	1000 - 10,000
Number of Robots	100
Number of Eggs	20%
Number of Free Riders	80%
Minimum organism size	2
Maximum organism size	10
Run length	100,000 time steps
Recruitment range	$\sim$ 3 robots (100 pixels)
Genome broadcast range	$\sim$ 3 robots (100 pixels)
Arena Size	64×64 robots (2048×2048 pixels)

III. Results and Analysis

Figure 1 shows the populations of two parameters sets where only the organism lifetime differs. The x axis shows the simulated time in time steps and the y axis shows the mean number of seeds, eggs and adult organisms, as well as the mean organism size of adult organisms. The first thing we notice is that the number of adult organisms when using a lifetime of $10{,}000$ quickly becomes much larger than when using a lifetime of 1000, this difference of 6.6 robots on average is significant: $p < 0.01$. This trend of a larger number of organisms when the organism lifetime is higher is reflected by all experiments.

In the experiments with many organisms the organism population starts fairly unstable with a big drop when the first generation of organisms dies, but the stability increases during the run to a point where it's very stable. The stability of the egg and seed populations seems to be reasonable as well and also increases over the course of the experiment, however they do not become as stable as the organism population. There seems to be an inherent periodicity in the egg and seed populations, which most likely the result of seeds that fail to build their organism which become an egg again. The mean organism size is also more stable in the second run, this can be attributed to the increase in number of adult organisms, but it is surprising how quickly this stabilises on an almost constant value of 3.5.

Lastly we notice that as the number of organisms rises, the number of eggs and seeds decreases. The effect of a larger organism population on the number

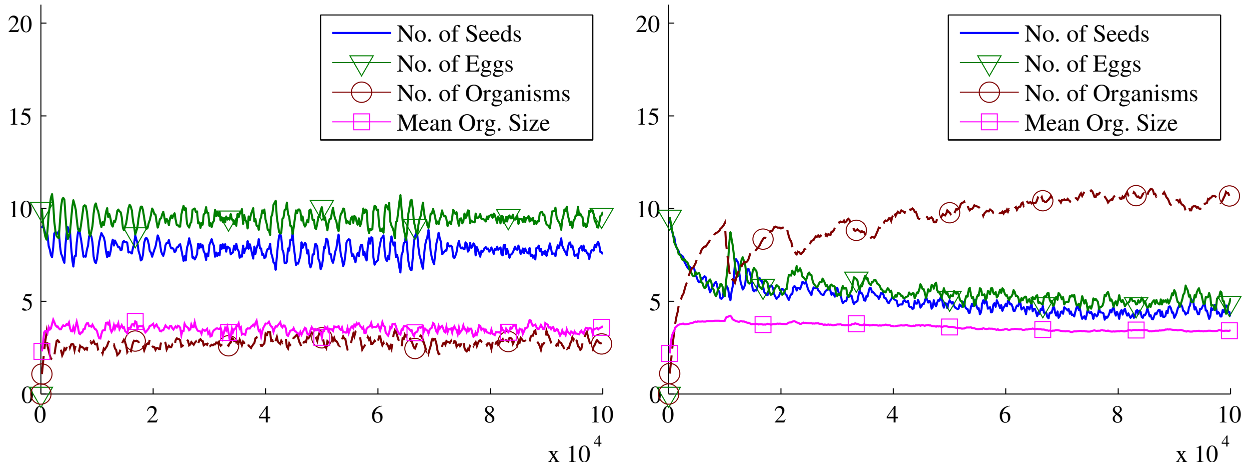

(a) Egg-, Seed- and Organism lifetime of 1000, 1000, 1000 (b) Egg-, Seed- and Organism lifetime of 1000, 1000, 10,000

Figure 1: Graphs showing a higher population for higher organism lifetimes

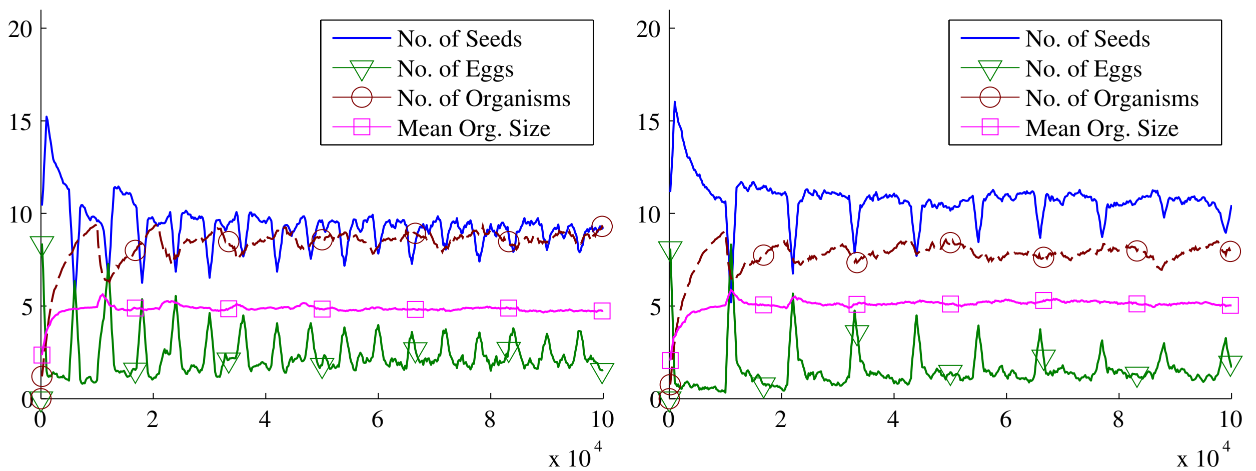

(a) Egg-, Seed- and Organism lifetime of 1000, 5000, 10,000 (b) Egg-, Seed- and Organism lifetime of 1000, 10,000, 10,000

Figure 2: Graphs showing a high periodicity for egg and seed populations and a high seed population

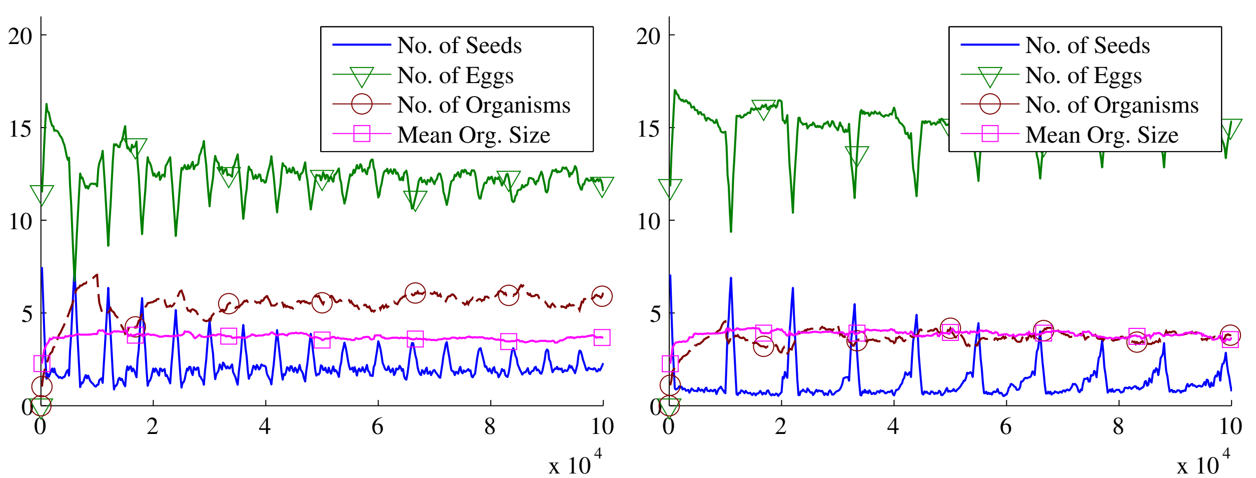

(a) Egg-, Seed- and Organism lifetime of 5000, 1000, 10,000 (b) Egg-, Seed- and Organism lifetime of 10,000, 1000, 10,000

Figure 3: Graphs showing a high egg population due to a higher egg lifetime

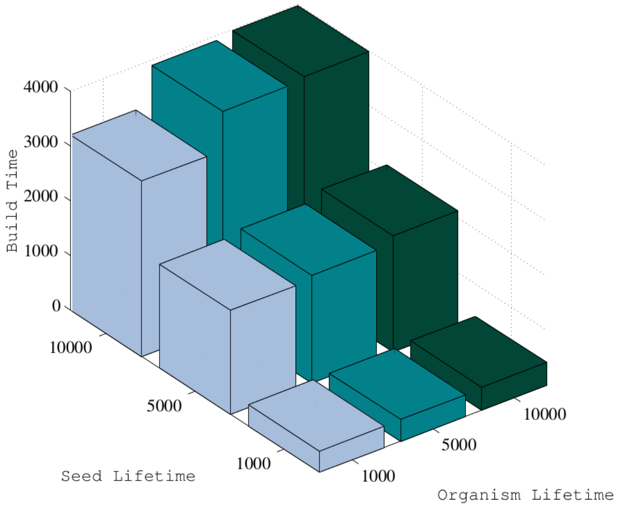

(a) Build Times for Egg lifetime of 1000

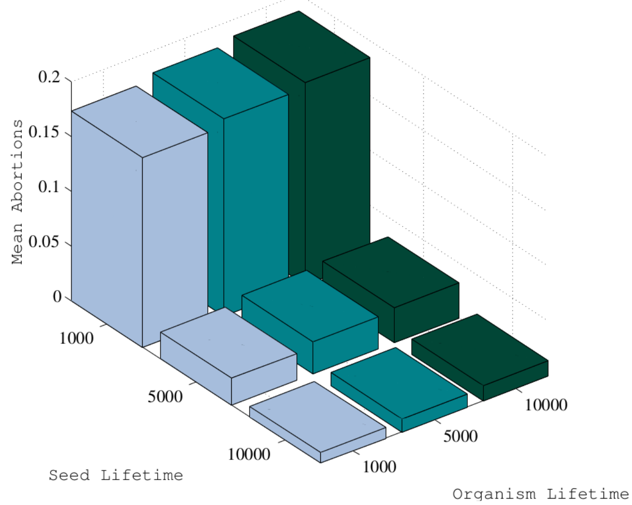

(b) Fraction of Abortions for Egg Lifetime of 1000

Figure 4: Comparison of abortion rates and build times, please note: the axis for seed lifetime is inverted in the abortion figure

of eggs and seeds results from the fact that more free modules are locked in organisms and thus not available to build organisms with. As eggs have a fixed lifetime this effect is very direct and its population is almost halved. For seeds this effect is less profound, the number of seeds decreases, however not as much as the egg population. The reason for this is that the lifetime of a seed is a *maximum* lifetime, and so varies from organism to organism, this is reflected in two measurements: the mean time to build an organism and the fraction of seeds that fail to build the organism and abort. Figure 4a shows the mean build time of organisms that completed and Figure 4b shows the mean fraction of seeds that aborted creating the organism due to a lack of time. The mean build time of organisms increases as the lifetime of organisms, and thus the number of organisms, increases. This trend is more visible for runs with a longer seed lifetime. Furthermore the mean fraction of abortions also increases as the organism lifetimes increases, although only slightly. Both these graphs confirm our hypothesis that a higher organism population makes it harder to build organisms, and thus the number of seeds is less affected by the number of organisms.

Figure 2 shows the populations of two parameter sets in which the seed lifetime differs. The first thing we notice is that the system becomes a lot less stable when compared to Figure 1. There is also a distinct periodicity in the number of eggs and seeds, even though we initialised the system to suppress periodicity. This periodicity of the egg and seed populations does seem to diminish over time, and although the populations do not stabilise completely in the course of these runs, we can assume it will become even more stable if we let the system run for longer.

Secondly the number of robots that are seeds is much higher when the lifetime of seeds is higher. Comparing the mean number of seeds between the three setups we find a difference of 4.3 robots between a seed lifetime of 1000 and 5000, which is significant ($p < 0.01$), secondly there is a difference of 1.4 robot between a seed lifetime of 5000 and $10{,}000$, which is also significant ($p < 0.01$). The increase in seed lifetime increases the mean build time for organisms by a lot as can be seen in Figure 4a, this explains increased size of seed population. The other effect of the longer seed lifetime is that the fraction of abortions decreases dramatically as seen in Figure 4b.

With more time to build organisms and fewer abortions one would expect to see more and larger organisms. Both Fig. 1 and 2 show the mean organism size and, although difficult to see on figures of this size, there is a significant difference ($p < 0.01$) in the mean size of organisms between a seed lifetime of 1000 and one of 5000: on average a difference of 1.23. There is also a significant difference in mean organism size between a seed lifetime of 5000 and $10{,}000$, however this is much smaller: 0.26. This means the influence of the seed lifetime on the organism size decreases as it gets larger, this is due to fewer available modules again. As the mean organism size in the system increases more modules are caught in building and adult organisms, therefore building larger organisms requires progressively more time to recruit the necessary modules.

The increased lifetime of seeds also influences the number of organisms, on average there are 0.98 fewer organisms when the seed lifetime is 5000 compared to when it is 1000 which is significant($p < 0.01$). The setup with a seed lifetime of $10{,}000$ 0.62 fewer organisms than the setup with $10{,}000$, this is significant only with $p < 0.05$. So, surprisingly a longer seed lifetime does not lead to *more* organisms, but to *fewer* organisms. We can only attribute this decrease in number of organisms to the increased size of organisms.

Moreover the increased number of seeds primarily results in a big drop in the number of eggs. The mean number of eggs is 3.3 lower with a seed lifetime of 5000 compared to a seed lifetime of 1000, and another 0.8 lower with a seed lifetime of $10{,}000$. The resulting amount of eggs is very detrimental to the use of this scenario for evolution; to ensure a good evolutionary system the organisms need to be able to spread their genomes effectively and with such a low and unstable number of eggs this becomes much harder.

Figure 3 shows the population dynamics of two parameter sets in which only the egg lifetime differs. Here we see a similar picture as in Fig. 2, the egg and seed populations are much less stable and show more periodicity. Here it is the eggs that are predominant, and the size of the egg population increases as the egg lifetime increases. The mean number of eggs is 6.8 higher with an egg lifetime of 5000 compared to an egg lifetime of 1000 and 2.7 higher with an egg lifetime of $10{,}000$ than an egg lifetime of 5000, both of which are significant differences ($p < 0.01$).

When comparing the difference in number of adult organisms between Fig. 1b , Fig. 2 and Fig. 3 you can see that the drop is much larger with a longer egg lifetime. On average the number of adult organism drops by 3.9 when increasing the egg lifetime from 1000 to 5000, and it drops by another 1.8 when increasing the egg lifetime from 5000 to $10{,}000$. This time the size of organisms is not influenced much, there is only a slight drop when moving to higher egg lifetimes. It is clear that a very high egg lifetime compared to the seed lifetime is very unhealthy for the system, as it has a detrimental effect on the stability of the populations and the number of organisms.

Again these population figures are not suitable for an evolutionary system, this time we have more than enough eggs to fertilise, however we lack a large population of organisms to ensure a healthy amount of reproductive competition.

Overall we can conclude that organism lifetime has the largest influence: increasing it always leads to more organisms. Egg lifetime is the next most influential. Decreasing it results a larger drop in number of organisms and increasing it leads to a larger increase in the number of eggs than does changing the seed lifetime. Thus, seed lifetime ranks third in terms of influence. All parameters do affect the population sizes for organisms, eggs and seeds to a certain extent, however. We see that having the egg lifetime and seed lifetime unequal to each other leads to unstable and periodic population sizes of eggs and seeds, although the periodicity may be reduced by choosing these values to be relatively prime.

With these insights we can construct a set of parameters for an 'ideal world' for our scenario, at least in terms of the parameters investigated. The egg and seed lifetimes should be equal, or at least close, to each other, and the organism lifetime should be an order of magnitude bigger than these two values. Of course having a very large organism lifetime leads to slower evolution of organism: there are simply fewer organisms per timespan. Furthermore, if the size of the organisms matters a higher seed lifetime increases the average organism size, although this will also reduce the number of organisms.

IV. Conclusion and Further Research

We have introduced a new scenario for the evolution of robotic body shapes where organisms roam an arena and fertilise eggs, which become seeds that form organisms by recruiting free modules. We have investigated three parameters of this scenario: egg lifetime, seed lifetime and organism lifetime, in particular in terms of their influence on the stability and size of the populations of organisms, eggs and seeds. We ran a total of 27 experiments, trying 3 different settings for each parameter.

The experiments show that the scenario reacts very strongly to the setting of organism lifetime, which is very influential for creating big populations of organisms and should be an order of magnitude higher than the egg and seed lifetimes to achieve large and stable numbers of organisms. The system also reacts quite strongly to the egg lifetime value: setting this higher than the seed lifetime leads to the undesirable situation of many eggs and few organisms. The seed lifetime is of least influence to the system, however it is still important to set it right in relation to the egg lifetime. This parameter is most influential of the three on the size of the organisms.

The egg lifetime and seed lifetime should consequently be chosen (almost) equal to each other. This promotes stable numbers of eggs and seeds and ensures sufficient eggs to allow evolution. Increasing seed lifetime leads to bigger organisms and fewer failed morphogenesis attempts, but it does so at the expense of the size of the organism population. Overall, the behaviour of the system with properly chosen parameter settings is promising and we are looking forward to using it in further research.

This work was a preliminary investigation to this scenario. Now that we have a better understanding of the influence of the egg, seed and organism lifetime parameters we can start using this scenario in investigations towards the evolution of robot shapes based on tasks and/or environment. We also wish to investigate the other parameters of this system to better understand it, such parameters include the distance at which genomes can be transmitted, distance at which a module can be recruited, and the number of free modules available.

ACKNOWLEDGEMENTS

This work was made possible by the European Union FET project SYMBRION. under grant agreement 216342. The authors would like to thank Quintin Saijoen, Mark Hoogendoorn and our partners in the SYMBRION consortium for many inspirational discussions on the topics presented here.

REFERENCES

[1] Bianco, R., Nolfi, S.: Toward open-ended evolutionary robotics: evolving elementary robotic units able to self-assemble and self-reproduce. Connection Science 4, 227–248 (2004)
[2] Bredeche, N., Montanier, J.M., Liu, W., Winfield, A.F.: Environment-driven distributed evolutionary adaptation in a population of autonomous robotic agents. Mathematical and Computer Modelling of Dynamical Systems 18(1), 101–129 (2012)
[3] Christensen, A., O'Grady, R., Dorigo, M.: Swarmorph-script: a language for arbitrary morphology generation in self-assembling robots. Swarm Intelligence 2(2), 143–165 (2008)
[4] Şahin, E.: Swarm robotics: From sources of inspiration to domains of application. Swarm Robotics pp. 10–20 (2005)
[5] Eiben, A.E., Elia, D., van Hemert, J.: Population dynamics and emerging features in AEGIS. In: *et al*, W.B. (ed.) Proceedings of the Genetic and Evolutionary Computation Conference. pp. 1257–1264. Morgan Kaufmann (1999)
[6] Eiben, A.E., Kernbach, S., Haasdijk, E.: Embodied artificial evolution – artificial evolutionary systems in the 21st century. Evolutionary Intelligence 5(4), 261–272 (2012), http://www.few.vu.nl/~ehaasdi/papers/EAE-manifesto.pdf
[7] Groß, R., Bonani, M., Mondada, F., Dorigo, M.: Autonomous self-assembly in swarm-bots. IEEE Transactions on Robotics 22, 1115–1130 (2006)
[8] Hosokawa, K., Tsujimori, T., Fujii, T., Kaetsu, H., Asama, H., Kuroda, Y., Endo, I.: Self-organizing collective robots with morphogenesis in a vertical plane. In: Robotics and Automation, 1998. Proceedings. 1998 IEEE International Conference on. vol. 4, pp. 2858 –2863 vol.4 (may 1998)
[9] Kernbach, S., Hamann, H., Stradner, J., Thenius, R., Schmickl, T., Crailsheim, K., Rossum, A.C.v., Sebag, M., Bredeche, N., Yao, Y., Baele, G., Peer, Y.V.d., Timmis, J., Mohktar, M., Tyrrell, A., Eiben, A.E., McKibbin, S.P., Liu, W., Winfield, A.F.T.: On adaptive self-organization in artificial robot organisms. In: Proceedings of the 2009 Computation World: Future Computing, Service Computation, Cognitive, Adaptive, Content, Patterns. pp. 33–43. COMPUTATIONWORLD '09, IEEE Computer Society, Washington, DC, USA (2009)
[10] Kutzer, M.D.M., Moses, M.S., Brown, C.Y., Scheidt, D.H., Chirikjian, G.S., Armand, M.: Design of a new independently-mobile reconfigurable modular robot. In: Robotics and Automation (ICRA), 2010 IEEE International Conference on. pp. 2758–2764. IEEE (2010)
[11] Liu, W., Winfield, A.: Autonomous morphogenesis in self-assembling robots using ir-based sensing and local communications. Swarm Intelligence pp. 107–118 (2010)
[12] Nolfi, S., Floreano, D.: Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines. MIT Press, Cambridge, MA (2000)
[13] O'Grady, R., Christensen, A., Dorigo, M.: Autonomous reconfiguration in a self-assembling multi-robot system. Ant Colony Optimization and Swarm Intelligence pp. 259–266 (2008)
[14] O'Grady, R., Christensen, A., Dorigo, M.: Swarmorph: multi-robot morphogenesis using directional self-assembly. Robotics, IEEE Transactions on 25(3), 738–743 (2009)
[15] Rasmussen, S., Knudsen, C., Feldberg, R., Hindsholm, M.: The coreworld: Emergence and evolution of cooperative structures in a computational chemistry. Physica D: Nonlinear Phenomena 42(1), 111–134 (1990)
[16] Ray, T.S.: Is it alive or is it GA? In: Belew, R., Booker, L. (eds.) Proceedings of the 4th International Conference on Genetic Algorithms. pp. 527–534. Morgan Kaufmann, San Francisco (1991)

[17] Sayama, H.: Robust morphogenesis of robotic swarms [application notes]. Computational Intelligence Magazine, IEEE 5(3), 43 –49 (aug 2010)

[18] Watson, R.A., Ficici, S.G., Pollack, J.B.: Embodied evolution: Distributing an evolutionary algorithm in a population of robots. Robotics and Autonomous Systems 39(1), 1–18 (April 2002), http://eprints.ecs.soton.ac.uk/10620/

[19] Wei, H., Cai, Y., Li, H., Li, D., Wang, T.: Sambot: A self-assembly modular robot for swarm robot. In: 2010 IEEE International Conference on Robotics and Automation (ICRA). pp. 66–71. IEEE (2010)

[20] Yim, M., Shen, W.M., Salemi, B., Rus, D., Moll, M., Lipson, H., Klavins, E., Chirikjian, G.S.: Modular self-reconfigurable robot systems [grand challenges of robotics]. Robotics & Automation Magazine, IEEE 14(1), 43–52 (2007)