

Open-ended On-board Evolutionary Robotics for Robot Swarms

Guy Baele^{*§}, Nicolas Bredeche[†], Evert Haasdijk[‡], Steven Maere[§], Nico Michiels[¶], Yves Van de Peer[§],
Thomas Schmickl^{||}, Christopher Schwarzer[¶], Ronald Thenius^{||}

^{*} Ghent University, Applied Mathematics & Computer Science, Belgium

[†] TAO, INRIA Saclay, Univ. Paris-Sud, France

[‡] Computational Intelligence Group, Amsterdam Free University, Netherlands

[§] Ghent University, Department of Molecular Genetics, Belgium

[¶] Animal Evolutionary Ecology Zoological Institute, University of Tuebingen, Germany

^{||} Department for Zoology, University of Graz, Austria

Abstract—The SYMBRION project stands at the crossroads of Artificial Life and Evolutionary Robotics: a swarm of real robots undergoes online evolution by exchanging information in a decentralized Evolutionary Robotics Scheme: the diffusion of each individual's genotype depends both on its ability to survive in an unknown environment as well as its ability to maximize mating opportunities during its lifetime, which suggests an implicit fitness. This paper presents early research and prospective ideas in the context of large-scale swarm robotics projects, focusing on the open-ended evolutionary approach in the SYMBRION project. One key issue of this work is to perform on-board evolution in a spatially distributed population of robots. A real-world experiment is also described which yields important considerations regarding open-ended evolution with real autonomous robots.

I. INTRODUCTION

Recent years have seen a growing interest from the Evolutionary Robotics community in robot swarms, both from fundamental and practical points of view: they provide a natural way to study evolutionary dynamics in a distributed set-up and address key issues regarding the increased robustness of a swarm of robots compared to one single robot with more features. Among the major ongoing projects in this field is the EU-funded project “SYMBRION”, which stands for “Symbiotic Evolutionary Robot Organisms” (see Fig. 1 for an illustration). This robot swarm project addresses both hardware, middleware and software issues, the latter including evolutionary approaches to robot behavior design.

The main focus of the SYMBRION project¹ is to investigate and develop novel principles of adaptation and evolution for symbiotic multi-robot organisms based on bio-inspired approaches and modern computing paradigms². Such robot organisms consist of super-large-scale swarms of robots that can dock with each other and symbiotically share energy and computational resources within a single artificial life-form. When it is advantageous to do so, these swarm

robots can dynamically aggregate into one or more symbiotic organisms and collectively interact with the physical world via a variety of sensors and actuators. The bio-inspired evolutionary paradigms, combined with robot embodiment and swarm-emergent phenomena, should enable the organisms to manage their own hardware and software organization autonomously. In this way, artificial robotic organisms become self-configuring, self-healing, self-optimizing and self-protecting from both hardware and software perspectives. Not only does this leads to extremely adaptive, evolvable and scalable robotic systems, it also enables robot organisms to reprogram themselves without human supervision and to develop new, previously unforeseen, functionality. In addition, different symbiotic organisms may co-evolve and cooperate with each other and with their environment.



Fig. 1. A model of aggregated swarm robots in the SYMBRION project.

There are other large-scale swarm robotics projects, such as the MIT swarm robot project [33] and EU funded SWARM-BOT [47], SWARMANOID [46] and I-SWARM [29] projects. These projects are representative of two complementary approaches to the design of swarm robot controllers. On the one hand, McLurkin's work on homogeneous swarm robots mainly focuses on designing ad hoc efficient and reliable controllers to address a wide range of swarm-related tasks (following a leader, chain formation, robot dispersion, etc.). On the other hand, the SWARM-BOT project addresses similar tasks with a different approach: that of automatically synthesizing robotic controllers through the use of an artificial

Contact author: nicolas.bredeche@lri.fr

¹The SYMBRION project is an EU funded FET started in January 2008 for a duration of five years under grant agreement 216342.

²A complete description of SYMBRION project is available here: www.symbion.eu

evolution algorithm. While actual validations on real robots are rare, some important results showed that the synthesized behaviors could be successfully transferred to a real world robot swarm for low-level tasks such as wall and hole avoidance by a group of physically connected robots. The ongoing SWARMANOID project extends these concepts to consider a heterogeneous swarm. Lastly, SYMBRION originates from research within the I-SWARM project [29] and especially from the open-source SWARMROBOT projects. In order to address the issues of learning and adaptation, SYMBRION investigates a number of complementary approaches including Machine Learning, off- and on-board Evolutionary Robotics. This paper focuses on the latter: the ability to provide self-adaptation through a continuously running evolutionary process that takes place as the robots perform their tasks. The key feature of this approach is that it provides an ever-adapting capability that does not require off-line simulation or central control. A positive side effect is that this enables researcher straight-forwardly to address the issue of the reality gap, often considered as one of the main deadlocks in Evolutionary Robotics [30].

The remainder of this paper states motivations and describes preliminary evolutionary models as well as hardware contingencies from preliminary experiments regarding open-ended on-board evolution within SYMBRION. In the next section, earlier works from the literature in Evolutionary Robotics and Artificial Life are reviewed and compared in the context of evolving populations of cooperative/competitive agents as well as SYMBRION originality and motivations with regards to previous works. Next, Sections III and IV provide a description of the seminal evolutionary operators (selection, replacement and variation) and representation issues, addressing genotype (what is evolved) and neuro-controller issues (how is it used to control robots). Section V describes an experimental set-up with a real swarm of homogeneous robots and provides preliminary results and insights regarding the on-line evolution of a swarm of robots in the real world. Finally, general considerations and further directions are discussed in section VI.

II. TOWARDS OPEN-ENDED EVOLUTIONARY ROBOTICS

A. Background

Evolutionary Robotics has been widely studied for over 20 years, starting with the work at EPFL [39] and at Sussex University [28] in the early 1990s. The first stages mainly focused on single autonomous systems, with a few notable exceptions studying co-evolutionary competitive processes [38] or cooperative processes [32], yet with (very) few individuals. However, recent projects focus on larger populations of interacting autonomous robots, be it robot swarms (SWARMBOT and SWARMANOID) or decentralized modular robots with multiple autonomous units (the M-tran modular robot project [26]). These projects rely on evolutionary-inspired methods to optimize a pre-defined criterion, such as locomotion or explicit coordination towards a fixed, supervisor-defined goal.

In parallel, in Artificial Life, there has been long-standing interest in self-replication as well as open-ended evolution, which has been simulated first in the work of Steen Rasmussen [41] and Thomas Ray [42]. This research focuses on the dynamics of evolution over time regarding an implicit criterion related to the ability of genes to spread within a population. Contrary to evolutionary robotics, it considers the dynamics of the whole population through locally interacting individuals by combining implicit selection with regards to survival capability and variation (such as mutations) to favor the emergence of fitter individuals. The key features in this set-up are:

- the ability to formulate implicitly some fitness function as a pressure towards survival (e.g., lifetime, ability to reproduce) from which derives the necessity to display more complex behaviors (exploration to find a mating partner, coordination for reproduction, etc.)
- the existence of an environment within which emergence of interactions and self-organization can be observed (i.e., considering interaction with both the world and other individuals).

Other works with comparable set-ups have studied, both on an individual as well as a collective basis, the emergence of complex behavior through cooperation and competition between clusters of individuals in the context of agent-based social simulation [20] or distributed virtual robotic agents moving in a two-dimensional simulated environment such as in AVIDA [1], Division blocks [45] and Polyworld [52], to name a few.

While Artificial Life researchers have pursued a so-called objective-free approach in contrast to the explicitly written objectives common in evolutionary robotics, there is growing interest from evolutionary robotics in the novelty that can emerge from open-ended evolution [7]. Broadly, this can be reduced to comparing two different approaches. One approach is to focus on exploring candidate solutions in a parallel fashion, favoring diffusion of useful traits for survival in a competitive/cooperative set-up. Another approach focuses on producing one specific optimized solution, whatever the structure of possible solutions, for which the problem of getting stuck in a local optimum –often due to lack of diversity– is a known and frequent problem. Indeed, the ability to generate novel behavior has recently been considered as a key feature to solve evolutionary robotics problems, either by directly favoring novelty search [31] or favoring diversity of comparably efficient individuals [37].

B. Brief outlook of evolutionary concepts in SYMBRION

In SYMBRION, our proposal is to rely on implicit pressure by binding the ability to spread one's genotype to both lifetime (i.e., individuals get more opportunities to spread genotype if they live longer) and mating behavior (i.e., individuals must find mates to reproduce, competition/cooperation may occur). Moreover, the environmental set-up must be carefully designed so as to favor emergence of complexity through a set of environmental constraints that act –similar to other

open-ended evolutionary systems— as evolutionary pressure, amongst others:

- constraints forcing implicit competition/cooperation between individuals (i.e., performance heterogeneity favors best performing individuals or group of individuals)
- locality of reproduction/replacement (i.e., diversity may be caused by spatial constraints)
- levels of complexity (i.e., the environment should be rich enough so as to avoid stalled equilibria, just as size was a continuous complexity ladder in Ray's [42] evolving programs).

SYMBRION offers a large-scale experimental set-up to study the emergence of behaviors in the context of a decentralized population of agents interacting locally. Moreover, the ability for robots to link physically to each other (e.g., to reach new places in the environment) offers opportunities for novel behavior and levels of coordination. Among the key issues to be studied are the necessary features for such a system with regards to important issues for evolutionary robotics: novelty, cooperative/competitive behaviors, self-organization and specialization in homogeneous vs. heterogeneous populations. Direct consequences for evolutionary robotics, such as the trade-off between implicit open-evolution-targeted fitnesses and directed pre-defined fitness will also be researched from the viewpoint of optimization as a discovery algorithm (i.e., the ability to provide sets of heterogeneous candidate solutions).

III. EVOLUTIONARY OPERATORS

A. General Scheme

Figure 2 shows a simplified illustration of the traditional Evolutionary Computation scheme. While the overall scheme for open-ended evolution in SYMBRION may differ in some aspects, the main operators (selection, replacement, variation) as well as the existence of an evaluation process resulting in some reward/fitness for each individual still exist and will be referred to in the following sections of this paper.




Fig. 2. Classic Evolutionary Computation Scheme

The main differences with the traditional scheme, which motivate the proposal of new operators to address the task at hand, are the following:

- evolution is an *asynchronous* process; each individual in the population may be “evaluated” at a different time, or even continuously (e.g., the ability to survive is a key advantage to diffuse one’s genotype);

- the evolutionary process is *distributed* over space (e.g., no central process may rank all individuals);
- as a result of the previous point, evaluation is *intrinsic* to the individual (e.g., there is no absolute ranking).

In the following section, selection and replacement will be described, as well as an intermediate “mating” operator that makes it possible to approximate to some extent the global behavior of a centralized evolutionary process. Variation operators, along with the related genotype representation issues, will be addressed in section IV.

B. Selection and Replacement Operators

Traditionally, selection and replacement operators apply selection pressure to the population of individuals, resulting in a trade-off between exploitation of highly ranked individuals and exploration of lower ranked but eventually more promising candidate solutions. This is usually achieved through a set of combined selection/replacement operators that range from fully deterministic ($\mu + \lambda$ -ES selection/replacement scheme) to probabilistic operators (tournament selection with low tournament size, which favors low ranked individuals) [16].

In a decentralized Evolutionary Algorithm the adaptive mechanisms have to be both **robust** and **scalable**: robust in the sense that the system has no single point of failure and can continue to function when robots are added or removed from the system; scalable in the sense that it can be used with large numbers of robots and in large environments without incurring undue overhead. These requirements lead to a design where robots adapt *autonomously* on the basis of only *local* information. This conflicts with most traditional evolutionary algorithms where some central authority evaluates, selects and replaces individuals. This crucial lack of a central authority implies that the agents must be able to gauge and compare their utility in some way. As quoted from [15]:

“The key element here is the locally executable selection. Crossover and mutation never involve many individuals, but selection in EAs usually requires a comparison among all individuals in the population.”

In other words, the operators that must differ from traditional evolutionary algorithms are the selection operators: mate selection instead of parent selection and a replacement strategy with only local control. The variation operators (crossover, mutation) need not be specific to embodied evolution.

For an evolutionary algorithm to converge to an acceptable solution, it is essential that better performing solutions produce more offspring than their lesser brethren. Therefore, a robot’s “fecundity” must relate to its performance either vis-à-vis the rest of the population or against some known maximum. Establishing the latter is often not possible or requires central control. To gauge one’s own fitness³ as well as that of possible mates, a robot can construct a sample

³‘Fitness’ in the sense of evolutionary algorithms: it denotes utility, not the actual rate of procreation as in the biological sense

from the robots it encounters (when moving around) or its neighbors (when in a fixed topology). In almost all cases this boils down to some variant of tournament selection as used in traditional evolutionary algorithms. Alternatively, one can estimate the population average and thus one's own relative performance through the gossiping algorithm as shown in [50], but even in this case, the pool of mates –if control is to remain local– must consist of a sample as described above.

In evolutionary robotics, there are surprisingly few forays into fully autonomous, localized evolution schemes. Most follow up –if sometimes only notionally– on Embodied Evolution, which was introduced by Watson, Ficici and Pollack [22], [49] and relies on broadcast of (mutated) genes at a rate proportional to the robot fitness. Some extensions of these original works were conducted either by introducing a *maturation period* [51], or by implementing time-sharing to cope with small populations [19],[18],[48]. Other work focuses on the competitive diffusion of genotype through comparing fitness value [43], spatially structured evolution strategies [25] and other local deterministic or tournament selection scheme using various genotypes and fitness values diffusion algorithms [11], [17], [3], [50].

In SYMBRION, two directions are considered. The first is inspired by the work on social learning done by [44]: agents volunteer as partners by sending ‘plumage’ containing ‘eggs’ at fixed intervals. Whenever a given fixed number of eggs has been received, the best is chosen for reproduction. To some extent, this emulates tournament selection [11], but with a sampling process that is biased by the very behavior of the robot (e.g., locality compromises uniform sampling). The authors show some interesting results regarding some parameters of the algorithm, such as some kind of elitism (in this context: accepting only eggs with higher fitness than one's own) may lead to slower convergence time because of resulting communication overhead.

Secondly, we propose an approach where robots impart their genetic material with a portion of their energy (which is independent from any explicit task realization criterion or battery level). In this set-up, the amount of energy imparted determines which candidate wins (either to fertilize an egg or to lay an egg in a ‘dead’ individual). Individuals die if their energy runs out and another fertilized robot may then implant its egg, with attached “dowry”. The dead individual remains open for implantation for a certain period, after which the egg with the highest dowry is used to reprogram the controller. The fraction of energy invested is determined genetically. This shares some strong similarities with bidding problems, and more specifically with sealed first-price auction [23], where the highest bidder gets the prize, but accepts to pay more than its counterparts. One key difference, however, lies in the fact that the bidding process is dependent from the exploratory behavior (i.e. ability to encounter mating partners) and results in optimizing the trade-off between exploration with environment constraints (i.e., the more energy conservative the longer the exploration) and reproduction constraints (i.e. the more energy dissipative, the higher the

probability to diffuse one's genotype). Section V shows some experiments with real robots using this set-up.

IV. REPRESENTATION ISSUES

A. Modular Genotype and Variation Operators

Designing a genome which immediately models a wide variation of evolutionary aspects at the nucleotide level and uses the properties of the genetic code to construct a protein layer, is inherently difficult. Different sections of DNA would then have to be modeled, corresponding to coding and non-coding genes. In cells, a gene is a portion of an organism's DNA which contains both coding sequences that determine what the gene does, and non-coding sequences that determine when and where the gene is active (i.e., expressed). When a gene is active, the coding and non-coding sequences are copied in a process called transcription, producing an RNA copy of the gene's information. This piece of RNA can then direct the synthesis of proteins via the genetic code.

Modeling changes and dependencies at the nucleotide level may quickly lead to attempts at modeling the complete path from changes of a single nucleotide to the corresponding consequences for e.g., the transcription process. In an initial phase, such interactions will not be present in the structure of the artificial genome (although they can be added in a later phase). Rather than starting from such a low-level scenario, the concept of genes and interactions between various genes is used as the starting point for the artificial genome. Some earlier works did show that such an approach is sounded even in the scope of Evolutionary Computation and should be explored further, both in a robotic context [14], [8] or in the general case [4].

Basic structure and properties: From a robotics point of view, it is interesting to think of the genome as a matrix decision structure because such a concept is easily translated into computer language. In such a framework, genes can be represented as small matrices, with each position in the matrix describing a certain property or characteristic of such a gene. A representation of the artificial genome can be seen in Figure 3, where the initial level at which the genome will be represented is the gene. Such genes will each have a current value, which will in turn be influenced by the values (i.e. the current characteristics) of other genes. It can be seen in Figure 3 that the value of the arbitrarily named “S” gene depends on the value of the “A”, “F”, “J”, ... genes. In other words, the value of S equals $23 + 0.4$ times the value of A - 0.2 times the value of F, ...

As a result, all genes can influence the expression of all other genes (both in a positive or negative way). This regulation can be set equal to zero at the beginning and is allowed to develop (evolutionarily) over time. Each regulatory relationship with another gene is also governed by its own specific mutation rate, which allows for fast evolving versus more conservative traits.

In biology, the linear dependencies depicted in Figure 3 are often oversimplifications and more complex dependencies are present, which use sigmoid functions [10]. Further, each

How S is influenced by genes A, F, J, ...						
S	A	F	J	.	.	.
value	23	$+0.4 \cdot A$	$-0.2 \cdot F$	$+0.01 \cdot J$		
mutation rate	0.01	0.10	0.07	0.00		
mutation step (can be prob. function)	0.003	0.01	0.06	0.01		

Fig. 3. Initial representation of several genes in an artificial genome, along with the interactions between the different genes, each gene's mutation rate and mutation step.

gene has a mutation rate, which can be altered by a certain mutation step. This does not have to be a simple increment or decrement of the mutation rate by a fixed value (i.e. the actual mutation step) but can take on more functions, such as a (normal) distribution with a mean and standard deviation.

This artificial genome representation can, in a first phase, include boolean operators (between the genes) and gene-specific fields that can be represented as a boolean variable. This could be used to turn a gene "on" or "off". Other genes may be allowed to depend on the status of this boolean variable, so changing the status can have a cascading effect on the other genes. For example, if the genes in the artificial genome are allowed to be regulated by transcription factors, these transcription factors could be able to give feedback to the genome and alter the state of the corresponding variable.

Advanced structure and properties: The basic structure of the artificial genome as well as its properties are initially fixed. However, incorporating the effects of certain evolutionary processes that have been shown to exist in nature requires an adaptation of the approach mentioned before. When evolutionary events, such as gene duplications (e.g., gene A is duplicated, leading to a new gene A*), are taken into account, the artificial genome becomes less static and can evolve new traits or characteristics.

Such changes need to be communicated to other parts of the robotic system, so that the actual robot can take advantage of this new gene with a currently unknown function and behavior. If the artificial genome aims to incorporate or mimic real world situations concerning gene duplication, the new gene A* will in most cases be deleterious. However, in certain cases (this could be simulated by allowing the new gene A* to remain in the artificial genome with a low probability) the new gene A* can be used to increase the capabilities of the robot so that the robot can acquire a new skill, giving it an evolutionary advantage over other robots.

Moreover, new skills may be acquired during the lifetime of a robot through adaptation of the controller, such as neural plasticity or learning mechanisms. Changes during lifetime are recorded in state vectors (see Figure 4) and may include various kinds of information such as internal states, immune-related information, hormone levels, learned effects, condition values, etc. These skills may eventually be reintegrated into the genome (following a Lamarckian scheme) or used as is without any further reintegration into

the genome (following a Baldwinian scheme) (see [40] for an introduction on both schemes).




Fig. 4. The role of the state vector. This figure shows the communication between this collection of phenotypes (i.e. the state vector) and the genome. The states in the state vector can affect the expression of genes and be affected in return.

However, it should be noted that there exists a wide range of possibilities between the exhaustive specification of the full controller and parameter tuning of an ad hoc hand-coded controller. In the first stage of the project, however, simple controllers with few parameters should be put to the test (e.g., implying weight optimization of a fixed topology multi-layer perceptron) while more promising, but difficult to tune, controllers will be left for further development (e.g., artificial ontogeny of neural network full topology and weights), as will be discussed in section IV-B.

B. Artificial Embryogeny for Neural Controllers

It has been shown that a combination of neural networks and evolutionary methods (termed neuroevolution) can perform better than traditional reinforcement learning methods in some specific domains, especially when the state space is vast and continuous and where no model of the environment may be available (i.e., where policy search is preferred). One advantage of neuroevolutionary methods is that the policy is represented using an artificial neural network (ANN), which is useful for learning tasks involving continuous (noisy) state variables [34]. The term embryogeny refers to the growth process which defines how a genotype maps onto a phenotype, possibly (however this is not mandatory) taken into account environmental feedback. Indeed, a cornerstone of biology is the distinction between inherited genetic code (the genotype which is inherited from the parents) and the corresponding organism (the phenotype, which results in the combination of the genotype intrinsic properties and of exogenous shaping, i.e., acquired features during lifetime).

Three different types of artificial embryogenies have been identified in evolutionary systems [6] so far: external, explicit and implicit. External means that the developmental process (i.e. the embryogeny) itself is not subjected to evolution but is hand-designed and defined globally and externally with respect to the genotypes. In explicit (evolved) embryogeny the developmental process itself is explicitly specified in the genotypes, and thus it is affected by the evolutionary process. Usually, the embryogeny is then represented in the genotype

as a tree-like structure following the paradigm of genetic programming. The third kind of embryogeny is implicit embryogeny, which comprises neither an external nor an explicit internal specification of the growth process. Instead, the embryogeny “emerges” implicitly from the interaction and activation patterns of the different genes. This kind of embryogeny has the strongest resemblance to the natural the evolutionary process.

Recent studies have explored such artificial ontogenic process, starting from a simple initial state, undergoing a sequence of changes (e.g., cell production) in order to reach a stable target organism state, both for evolving robotic morphologies [8] or in simpler set-up closer to continuous-state cellular automata [36], [21], [24], [12], focusing on scalability and robustness issues.

Other studies focused on the dynamic linking of cells during the developmental phase [27], [2], whereby both the direction of growth as well as the linking probabilities between nodes depend on inherited genetic information and on the modulation of axon growth induced by gradients of virtual chemicals (morphogenes). Within the SYMBRION project, we plan to combine these concepts and we plan additionally to simulate cell movement, cell death and cell differentiation. These concepts were successfully demonstrated in [5], [36] and are currently being investigated in the scope of this project [12], [13]. In this scope, the dynamics and the spatial organization of these developmental processes depend on information stored in the genome of the robot, which is shaped by an evolutionary process.

V. PRELIMINARY EXPERIMENTS AND DISCUSSION

In this section, a first experiment in on-board on-line evolution is described. As the SYMBRION robot prototypes are currently still under conception, all experiments were conducted with a set of 40 Jasmine III robots available from the swarmrobot project that preceded SYMBRION [29], see figure 5.



Fig. 5. An experiment run with about 40 robots. Wall sections in the field prevent that robots uniformly follow the walls. Gender and life status are indicated with colored LEDs.

The following experiment takes advantage of most, but not all, of the previously described concepts: open-ended evolution is achieved by an implicit fitness based solely on diffusion of the genotype (see section II-A) and relies on a selection/replacement operator that is based on the amount of energy a robot is prepared to invest so as to diffuse its genotype, as described in section III-B.

The next sub-section describes the experimental set-up as well as some results obtained with two different population. Next, some considerations about open-ended evolution are discussed as well as some insights regarding future experiments.

A. Experiment: Competition of different sexual populations

A concrete implementation of open-ended evolution has been developed on Jasmine III robots available at the University of Stuttgart for modeling biologically inspired sexual behavior. In this set-up, the motivation is to study whether a robot population with split functions for reproduction, like in males and females has advantages over a population where each robot has the same capabilities. The underlying principles state that evolution should be performed based on a local exchange of genomes between nearby robots, where each robot carries an active genome that regulates part of its behavior and impact its ability to survive.

In the scope of the current experimental set-up, each robot wanders around the environment according to a simple random exploration behavior biased towards forward movement. Each robot is endowed with an internal energy level that is full at birth and linearly decreases over time. Once energy is depleted, the robot is considered dead and is free for insemination by another robot's genotype. In this basic experiment, the genotype contains two evolving traits: the amounts of energy spent during mating and creating offspring. As a consequence, the amount of energy the robot is willing to invest to win a bidding process for inseminating dead robots directly impacts the reproduction ability of the individual. As described in section III-B, this means that the higher energy spent, the higher the probability to inseminate a dead robot successfully, as each robot competes with all other robots. Thus, there is a trade-off between few, but reliable, mating opportunities and many, but doubtful, mating opportunities.

The genetic exchange in a population is close to panmixia and the only hindrance are other robots that form a dynamic, spatial obstacle. Two sexual modes have been implemented. In one mode –gonochorism– sexual functions are split into two genders, male and female; only females can create offspring after receiving a genome from a male. In the other mode –hermaphrodisism–, each member of the population can perform both functions simultaneously. Apart from this difference both modes behave identically. These two modes of reproduction have different advantages and disadvantages depending on the environment. Biological models predict that, for example, hermaphrodisism is expected to be more efficient than gonochorism in a situation that permits only a low number of matings per life time [35].

The current experiment is defined as follows: the robot swarm is initially split into two equally sized populations of hermaphrodites and gonochorists. The two populations do not interact nor mate, however they do compete for dead robots to create offspring. Both sides can attempt to revive the same dead robot during a 6 second time span, after which the attempt with the highest energy spent (“dowry”) recruits

this robot as offspring. The population with the more efficient strategy will gradually increase in numbers because they win more revive attempts. Thus, a population can be overcome by the other over time and a run ends when only members of one population remain, unless there is a stable coexistence.

This experiment is performed over a range of different arena sizes to find out how the swarm robot density and thus the mating rate affects the efficiency of the two reproductive strategies. Robots are placed randomly in the arena shown in figure 5 and interact with one another whenever they come into communication range of the infrared equipment (about 6 cm). Each robot signals its status via an infrared broadcast, which is emitted frequently (about 100 times per second). Thus, robots can virtually perceive other robots despite their marginal sensory capabilities by reading their broadcasts. The maximum lifetime is set to 100 seconds, which is enough to keep the population from going extinct and leads to a large number of generations in the one hour run time of the robots' battery.

The runs were conducted in two arenas of different sizes and with two population sizes (20 and 40 robots) to study the influence of population density. Results were compiled from 38 independent runs, lasting a maximum of 8 minutes, to provide reliable figures. Hermaphrodites recruit most or all the robots within 8 minutes (from 55% of the runs in the large environment with a small population to 100% of the runs in the small environment with large population). This is not a surprise in this context as pressure selection is driven by the number of opportunities to mate rather than the acquisition of specific behavior. However, it should be noted that sexual selection may boost the gonochoristic reproductive strategy and generally increase the rate of evolution and adaptation. Indeed, theories from biology suggest that a key feature of gonochorism is a more efficient sexual selection[35] and may benefit the population in difficult environments where adaptation is strongly required. On-going research is currently investigating this in set-ups where energy gain is included in foraging behavior and/or as a reward for performing different kind of tasks. As producing offspring is costly in the current set-up, the total energy given at birth provides an evolutionary advantage and can be straightforwardly related to the achievement of user-defined tasks.

B. Discussion on Technical Considerations

As a complementary result of these experiments, it is apparent that some aspects of the hardware design of the robot are a major influence to the quality and effectiveness for this kind of evolutionary process such as (1) communication (robot identification and information exchange) and (2) motor capability (speed and agility). The good characteristics of Jasmine III in this aspect⁴ allowed a high rate of encountering different robots. However, some hardware aspects, such as the shape of the chassis, formed a very unexpected influence as robot crowding emerged from time to time due to collisions and the chassis becoming stuck to each other, literally

⁴Velocity is about 0.5 m/s, time for a full turn is less than 1.0 s.

forming death traps. Finally, considerations on the swarm size are highly dependent on the problem at hand, but a critical mass is required for reliable results. In the current set-up, the number of robots used is a critical issue - while no statistically significant results may be provided at this point, preliminary experiments showed that at least 40 robots were needed to overcome spontaneous extinction in the current setup.

VI. CONCLUSIONS AND PERSPECTIVES

This paper presents on-going research in the scope of on-board open-ended evolutionary robotics in the European IP SYMBRION project. An overview of the Evolutionary Approach scheme is provided, considering both evolutionary operators and preliminary considerations on representation issues regarding neuro-controllers. A simple, yet real-world experiment was described using real robots with on-board adaptation thanks to artificial evolution. The main contribution of this paper is two-fold as it provides a description of preliminary directions of a long-term large-scale swarm robot project and gives some concrete insights regarding real-world on-board open-ended evolution with many (up to 40) real robots.

Perspectives in the short term include extension of this work, in particular by considering the impact of selection pressure towards diversity of the population and its relations with the exploration/exploitation trade-off regarding adaptation as well as representation issues, with a special emphasis on artificial ontogeny for growing neural controllers. In the long term, the SYMBRION robot modules actually under construction are due to be able to move around and to connect with one another to form complex shapes such as snake or spider morphologies. This should bridge the gap between swarm and modular robotics. Future work will focus on the self-organization in this forthcoming set-up.

ACKNOWLEDGEMENTS

The SYMBRION project is funded by the European Commission through the FET Proactive Initiative: PERVASIVE ADAPTATION, the grant agreement no. 216342. Additionally, we want to thank all members of the SYMBRION and REPLICATOR⁵ projects for fruitful discussions.

REFERENCES

- [1] C. Adami and T. C. Brown. Evolutionary learning in the 2d artificial life system avida. In *Proceedings of Artificial Life IV*, page 377381, 1994.
- [2] C. Adami, N. Chaumont, J. Edlund, and A. Hintze. Topological properties of evolved robot brains (abstract). In S. Bullock, J. Noble, R. Watson, and M. A. Bedau, editors, *Artificial Life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems*, page 742. MIT Press, Cambridge, MA, 2008.
- [3] E. Alba and M. Tomassini. Parallelism and evolutionary algorithms. *IEEE Transactions on Evolutionary Computation*, 6(5):443–462, October 2002.
- [4] W. Banzhaf. On the dynamics of an artificial regulatory network. In *7th conference on Artificial Life*, pages 217–227, 2003.

⁵European IP project, funded within the work programm "Cognitive Systems, Interaction, Robotics" under the grant agreement no. 216240.

- [5] D. Basanta, M. Miodownik, and B. Baum. The evolution of robust development and homeostasis in artificial organisms. *PLoS Comput Biol*, 4(3):e1000030, Mrz 2008.
- [6] P. Bentley and S. Kumar. Three ways to grow designs: A comparison of embryogenies for an evolutionary design problem. In *Genetic and Evolutionary Computation Conference*, pages 35–43. Morgan Kaufmann, 1999.
- [7] R. Bianco and S. Nolfi. Toward open-ended evolutionary robotics: evolving elementary robotic units able to self-assemble and self-reproduce. *Connection Sciences*, 16:227–248, 2004.
- [8] J. C. Bongard and R. Pfeifer. Evolving complete agents using artificial ontogeny. In *In Morpho-functional Machines: The New Species (Designing Embodied Intelligence)*, pages 237–258. Springer-Verlag, 2003.
- [9] P. A. N. Bosman, T. Yu, and A. Ekárt, editors. *GECCO '07: Proceedings of the 2007 GECCO conference on Genetic and evolutionary computation*, New York, NY, USA, 2007. ACM.
- [10] A. Climescu-Haulica and M. Quirk. A stochastic differential equation model for transcriptional regulatory networks. *BMC Bioinformatics*, 8, 2007.
- [11] K. De Jong and J. Sarma. On decentralizing selection algorithms. In L. Eshelman, editor, *Proceedings of the Sixth International Conference on Genetic Algorithms*, pages 17–23, San Francisco, CA, 1995. Morgan Kaufmann.
- [12] A. Devert, N. Bredeche, and M. Schoenauer. Robust multi-cellular developmental design. In D. Thierens et al., editor, *GECCO'07*. ACM Press, 2007. Preprint available at <http://hal.inria.fr/inria-00145336/en/>.
- [13] A. Devert, N. Bredeche, and M. Schoenauer. Unsupervised Learning of Echo State Networks: A case study in Artificial Embryogeny. In N. Montmarché et al., editor, *Artificial Evolution, Selected Papers*, volume 4926/2008 of *Lecture Notes in Computer Science*, pages 278–290, 2007.
- [14] P. Durr, C. Mattiussi, and D. Floreano. Neuroevolution with Analog Genetic Encoding. In *Parallel Problem Solving from Nature - PPSN IX*, volume 9 of *Lecture Notes in Computer Science*, pages 671–680, 2006.
- [15] A. E. Eiben, M. Schoenauer, J. L. J. Laredo, P. A. Castillo, A. M. Mora, and J. J. Merelo. Exploring selection mechanisms for an agent-based distributed evolutionary algorithm. In Bosman et al. [9], pages 2801–2808.
- [16] A. E. Eiben and J. Smith. *Introduction to Evolutionary Computing*. Springer-Verlag, London, 2003.
- [17] S. E. Eklund. A massively parallel architecture for distributed genetic algorithms. *Parallel Computing*, 30(5 – 6):647 – 676, May – June 2004.
- [18] S. Elfving. *Embodied Evolution of Learning Ability*. PhD thesis, KTH School of Computer Science and Communication, SE-100 44 Stockholm, Sweden, November 2007.
- [19] S. Elfving, E. Uchibe, K. Doya, and H. Christensen. Biologically inspired embodied evolution of survival. In Z. Michalewicz and R. G. Reynolds, editors, *Proceedings of the 2008 IEEE Congress on Evolutionary Computation IEEE Congress on Evolutionary Computation*, volume 3, pages 2210–2216, Hong Kong, June 2008. IEEE Press.
- [20] J. Epstein and R. Axtell. *Growing Artificial Societies: social science from the bottom up*. MIT Press, 1996.
- [21] D. Federici. Increasing evolvability for developmental programs. In J. Miller, editor, *Workshop on Regeneration and Learning in Developmental Systems*, *WORLDS 2004*, 2004.
- [22] S. Ficici, R. Watson, and J. Pollack. Embodied evolution: A response to challenges in evolutionary robotics. In J. L. Wyatt and J. Demiris, editors, *Proceedings of the Eighth European Workshop on Learning Robots*, pages 14–22, 1999.
- [23] M. Flavio and M. Paulo. *An introduction to auction theory*. Oxford University Press, 2005.
- [24] T. G. W. Gordon and P. J. Bentley. Bias and scalability in evolutionary development. In *GECCO '05*, pages 83–90. ACM Press, 2005.
- [25] M. Gorges-Schleuter. A comparative study of global and local selection in evolution strategies. In A. Eiben, T. Th. Bäck, M. Schoenauer, and H.-P. Schwefel, editors, *Parallel Problem Solving from Nature – PPSN V*, volume Volume 1498/1998 of *Lecture Notes in Computer Science*, pages 367–377, Berlin / Heidelberg, 1998. Springer-Verlag.
- [26] H. K. H., E. Yoshida, K. Tomita, K. A., S. Muratab, and S. Kokajia. Self-reconfigurable m-tran structures and their walker generation. *Robotics and Autonomous Systems*, 54(2):142–149, 2006.
- [27] A. N. Hampton and C. Adami. Evolution of robust developmental neural networks, 2004.
- [28] I. Harvey, P. Husbands, D. Cliff, A. Thompson, and N. Jakobi. Evolutionary robotics: the Sussex approach. *Robotics and Autonomous Systems*, 20:205–224, 1997.
- [29] www.i-swarm.org.
- [30] N. Jakobi. Minimal simulations for evolutionary robotics. Technical report, University of Sussex, 1998.
- [31] J. Lehman and K. O. Stanley. Exploiting open-endedness to solve problems through the search for novelty. In *Proceedings of Artificial Life XI*, 2008.
- [32] A. Martinoli and F. Mondada. Collective and cooperative group behaviours: Biologically inspired experiments in robotics. In *Proceedings of the Fourth International Symposium on Experimental Robotics (1995)*, 1995.
- [33] J. McLurkin. Dynamic task assignment in robot swarms. In *Proceedings of Robotics: Science and Systems*, pages 2007–214838, 2005.
- [34] J. H. Metzen, M. Edgington, Y. Kassahun, and F. Kirchner. Evolving neural networks for online reinforcement learning. In *Parallel Problem Solving from Nature*, pages 518–527, 2008.
- [35] N. K. Michiels and J. M. Koene. Sexual selection favors harmful mating in hermaphrodites more than in gonochorists. *Integrative and Comparative Biology*, 46(4):473–480, 2006.
- [36] J. F. Miller and W. Banzhaf. Evolving the program for a cell: from french flags to boolean circuits. In S. Kumar and P. J. Bentley, editors, *On Growth, Form and Computers*. Academic Press, 2003.
- [37] J. Mouret and S. Doncieux. Incremental evolution of animats behaviors as a multi-objective optimization. In *Proceedings of the International Conference on Simulation of Adaptive Behavior (SAB)*, 2008.
- [38] S. Nolfi and D. Floreano. Coevolving predator and prey robots: Do arms races arise in artificial evolution? *Artificial Life*, 4:311–335, 1998.
- [39] S. Nolfi and D. Floreano. *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. The MIT Press, Cambridge, MA. / London, 2000.
- [40] S. Nolfi, S. Nolfi, J. L. Elman, J. L. Elman, D. Parisi, and D. Parisi. Learning and evolution in neural networks. *Adaptive Behavior*, 3:5–28, 1994.
- [41] S. Rasmussen, K. Carsten, P. Feldberg, and M. Hindsholm. Is it alive, or is it ga? *Physica D*, 42:111–134, 1990.
- [42] T. Ray. Is it alive, or is it ga? In *Proceedings of International Conference on Genetic Algorithms*, 1991.
- [43] E. D. V. Simões and K. R. Dimond. Embedding a distributed evolutionary system into population of autonomous mobile robots. In *Proceedings of the 2001 IEEE Systems, Man, and Cybernetics Conference*, 2001.
- [44] R. E. Smith, C. Bonacina, P. Kearney, and W. Merlat. Embodiment of Evolutionary Computation in General Agents. *Evolutionary Computation*, 8(4):475–493, 2000.
- [45] L. Spector, J. Klein, and M. Feinstein. Division blocks and the open-ended evolution of development, form, and behavior. In *Proceedings of GECCO*, 2007.
- [46] www.swarmanoid.org.
- [47] www.swarm-bots.org.
- [48] Y. Usui and T. Arita. Situated and embodied evolution in collective evolutionary robotics. In *Proceedings of the 8th International Symposium on Artificial Life and Robotics*, pages 212–215, 2003.
- [49] R. A. Watson, S. G. Ficici, and J. B. Pollack. Embodied evolution: Distributing an evolutionary algorithm in a population of robots. *Robotics and Autonomous Systems*, 39(1):1–18, April 2002.
- [50] W. R. M. U. K. Wickramasinghe, M. van Steen, and A. E. Eiben. Peer-to-peer evolutionary algorithms with adaptive autonomous selection. In Bosman et al. [9], pages 1460–1467.
- [51] S. Wischmann, K. Stamm, and F. Wörgötter. Embodied evolution and learning: The neglected timing of maturation. In *Advances in Artificial Life: 9th European Conference on Artificial Life*, LNAI, in press. Springer-Verlag, in press.
- [52] L. S. Yaeger, V. Griffith, and O. Sporns. Passive and driven trends in the evolution of complexity. In *Proceedings of Artificial Life XI*, 2008.