

# Evolving Team Compositions by Agent Swapping

Paweł Lichocki, *Member, IEEE*, Steffen Wischmann, Laurent Keller, and Dario Floreano, *Senior Member, IEEE*

**Abstract**—Optimizing collective behavior in multiagent systems requires algorithms to find not only appropriate individual behaviors but also a suitable composition of agents within a team. Over the last two decades, evolutionary methods have emerged as a promising approach for the design of agents and their compositions into teams. The choice of a crossover operator that facilitates the evolution of optimal team composition is recognized to be crucial, but so far, it has never been thoroughly quantified. Here, we highlight the limitations of two different crossover operators that exchange entire agents between teams: restricted agent swapping (RAS) that exchanges only corresponding agents between teams and free agent swapping (FAS) that allows an arbitrary exchange of agents. Our results show that RAS suffers from premature convergence, whereas FAS entails insufficient convergence. Consequently, in both cases, the exploration and exploitation aspects of the evolutionary algorithm are not well balanced resulting in the evolution of suboptimal team compositions. To overcome this problem, we propose combining the two methods. Our approach first applies FAS to explore the search space and then RAS to exploit it. This mixed approach is a much more efficient strategy for the evolution of team compositions compared to either strategy on its own. Our results suggest that such a mixed agent-swapping algorithm should always be preferred whenever the optimal composition of individuals in a multiagent system is unknown.

**Index Terms**—Cooperation, crossover, evolutionary computation, multiagent systems, team composition, team optimization.

## I. INTRODUCTION

THE OPTIMIZATION of collective behavior displayed by teams of agents plays a crucial role in an increasing number of applications [1], [2], spanning from software agents [3]–[5] to robotics [6]–[8]. Evolutionary computation has been advocated as an effective and promising strategy in this domain [9], [10]. An important question that arises has to do with the composition of the teams of agents. All agents from one team may either use the same control algorithm (genetically homogenous teams) or employ different ones

(genetically heterogenous teams) [11]–[13]. Evolving homogenous teams does not differ conceptually from evolving single agents because in both cases, only one control algorithm is discovered [14]. In contrast, with heterogeneous teams, a set of distinct control algorithms must be optimized simultaneously. Consequently, the challenge is to find not only the optimal agents but also the optimal composition of agents within a team [13]. Heterogenous teams are of growing interest in the evolutionary community because they are expected to perform better than homogeneous teams in problems that require task specialization [12], [13], [15].

In heterogeneous team evolution, two genetic encodings can be used: 1) individual encoding, where a genome represents one agent, and 2) team encoding, where a genome represents a whole team. Consequently, there are differences in the evolutionary algorithm used with individual and team encoding [Fig. 1]. With individual encoding, one must decide on a method of grouping agents into temporary teams for the purpose of performance evaluation [Fig. 1(a)]. With team encoding, one must choose a crossover operator that exchanges genetic material not only between agents but also between teams (i.e., swap agents) [Fig. 1(b)]. Both issues are facets of the same challenge: how to (re)compose agents in teams, in order to facilitate the evolutionary search. This question has already been studied for individual encoding [Table I, (A)–(C)] but scarcely addressed for team encoding, although the body of work using team encoding is rich [Table I, (D)–(G)].

In team encoding, a single genotype encodes the entire team, which makes it decomposable into parts corresponding to the agents. Thus, a crossover operator may exchange genetic material on two levels [14]. First, the crossover can recombine the genetic material between agents from the parenting teams. We refer to this process as agent recombination [Table I, (D), (E) and Fig. 2 (top row)]. Second, the crossover can swap entire agents between the parenting teams. We refer to this process as agent swapping [Table I, (F), (G) and Fig. 2 (bottom row)]. In contrast to agent recombination, agent swapping does not exchange genetic material between the agents. Consequently, the purpose of agent recombination is to discover good agents and the purpose of agent swapping is to discover good team compositions. In addition, one may consider a team to be an ordered sequence of agents. In such a case, the crossover may be restricted to act only on agents on corresponding positions in the parenting teams [Table I, (D) and (F) and Fig. 2 (left)]. Or, it may be free to act on any agents from the parenting teams [Table I, (E) and (G) and Fig. 2 (right)].

In contrast to agent recombination [16]–[18], no attempts have been made to quantify the efficiency of agent swapping in the evolution of teams [12]. In particular, it has not been

Manuscript received October 21, 2011; revised February 3, 2012; accepted March 4, 2012. Date of publication March 19, 2012; date of current version March 27, 2013. This work was supported by the Swiss National Science Foundation under Grant K-23K0-117914. The review of this paper was coordinated by H. Abbass.

P. Lichocki and D. Floreano are with the Laboratory of Intelligent Systems, École Polytechnique Fédérale de Lausanne, Lausanne 1015, Switzerland (e-mail: pawel.lichocki@epfl.ch; dario.floreano@epfl.ch).

S. Wischmann is with the Laboratory of Intelligent Systems, École Polytechnique Fédérale de Lausanne, Lausanne 1015, Switzerland, and also with the Department of Ecology and Evolution, Biophore, University of Lausanne, Lausanne 1015, Switzerland (e-mail: steffen.wischmann@epfl.ch).

L. Keller is with the Department of Ecology and Evolution, Biophore, University of Lausanne, Lausanne 1015, Switzerland (e-mail: laurent.keller@unil.ch).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TEVC.2012.2191292

TABLE I  
SAMPLE OF APPROACHES IN THE EVOLUTION OF HETEROGENEOUS  
TEAMS USING INDIVIDUAL ENCODING AND TEAM ENCODING

Individual Encoding	
Agent Grouping	References
(A) One team	[70]–[80]
(B) Many teams	[13], [18], [24]–[27], [81], [82]
(C) Subpopulations	[11], [25], [83]–[87]
Team Encoding	
Crossover	References
(D) Restricted agent recombination	[11], [14]–[18], [20]–[23], [27], [32], [70], [88]
(E) Free agent recombination	[16]–[18], [20], [89]
(F) RAS	[22], [23]
(G) FAS	–

tested if and why using agent swapping leads or does not lead to the evolution of optimal team compositions. Here, we experimentally compare restricted agent swapping (RAS) and free agent swapping (FAS) in a problem focused on finding the optimal team composition. We consider multiple agents that need to divide the labor in order to achieve top performance, i.e., the optimal team is composed of distinct groups of genetically identical agents. We focus on team encoding, which assumes team level of selection [Fig. 1] that has been advocated as an efficient strategy in the optimization of teams [13]. Consequently, agent interactions and their impact on individual selection pressures are out of scope of this paper. Also, we focus on the evolution of team compositions and not on the evolution of agents themselves. Thus, we mainly consider large teams of agents having small genomes, i.e., a valid single agent is relatively easy to evolve. We further elaborate on these assumptions in Section V.

We highlight the limitations of both RAS and FAS and explain the conditions under which they fail to evolve teams displaying the optimal composition. These limitations are opposite to each other. With RAS, the evolutionary algorithm suffers from premature convergence of the population, whereas with FAS, it suffers from insufficient convergence of the population. Consequently, in both cases, the exploration and exploitation aspects of the evolutionary algorithm are not well balanced. To overcome this problem, we propose to combine the two methods, i.e., first use FAS to explore the search space and then use RAS to exploit it. This mixed approach proves to be a more efficient strategy in the evolution of team compositions than RAS or FAS alone. Finally, we also validated RAS and FAS on a problem of optimizing decentralized controllers for task allocation and discussed our results in the context of other real-life applications.

## II. BACKGROUND

The differentiation into individual and team encoding resembles an old discussion on evolving rule-based systems [19]. In the approach taken by De Jong, dubbed “the Pitt

approach,” a single individual encoded the entire rule set. In contrast, in the approach taken by Holland, dubbed “the Michigan approach,” a single individual encoded just a single decision rule and the entire population corresponded to the rule set (see [19] and references therein).

The approaches using individual encoding can be classified into three categories, according to how the agents are grouped into teams for the purpose of performance evaluation. With the “one team” approach, all agents from the population are evaluated together, i.e., they compose one team [Table I, (A)]. This method is often used with a continuously updated gene-pool, i.e., steady-state evolution. With the “many teams” approach, the agents are randomly grouped into many teams [Table I, (B)]. This method is often used with separate gene-pools for subsequent generations, i.e., generational evolution. With the “subpopulations” approach, there are separate subpopulations of agents [Table I, (C)]. An individual is evaluated by teaming it up with individuals from other subpopulations. This method is known as cooperative coevolution.

In individual encoding, the teams are created *ad hoc* for the purpose of performance evaluation. Consequently, there are no genetic operators applied at the team level. In contrast, in team encoding, we distinguish four qualitatively different types of crossover operators that exchange genetic material between the teams [Fig. 2]. The crossover either *recombines* the agents from the parenting teams [Table I, (D) and (E)] or it *swaps* entire agents between the parenting teams [Table I, (F) and (G)]. In addition, the crossover is either *restricted* to act only on agents on corresponding positions in the parenting teams [Table I, (D) and (F)] or it is *free* to act on any agents from the parenting teams [Table I, (E) and (G)]. Note that agent swapping is in fact a special case of agent recombination, where the recombination points are always chosen at the beginning (or at the end) of the two genotypes’ parts that encode the parenting agents.

The concept of restricted (also called fixed in [18]) and free (also called unfixed in [18] and interpositional in [20]) agent recombination was introduced independently by Luke and Spector [16] and by Haynes and Sen [21]. Agent recombination was called inner crossover in [14]. RAS, called team transformation in [22], was introduced by Luke *et al.* [23] and by Andre and Teller [22]. FAS has not been studied directly yet [Table I, (G)]. Note that evolutionary algorithms with individual encoding, random grouping of agents into many teams, and global reward, yield high resemblance to team encoding with FAS [Fig. 1]. This includes some of the work referenced in [Table I, (B)], i.e., [13], [18], [24]–[27].

Haynes and Sen [21] noted that RAS may be implemented as *uniform* crossover (i.e., swapping bits on corresponding positions between two genotypes with some probability), with the difference that instead of bits it swaps entire agents. Here, we analogously note that FAS may be implemented as *shuffle-uniform* crossover (i.e., swapping bits on shuffled positions between two genotypes with some probability), with the difference that instead of bits, it swaps entire agents. The family of uniform crossovers was introduced by Syswerda [28] and analytically studied by Eshelman *et al.* [29] and

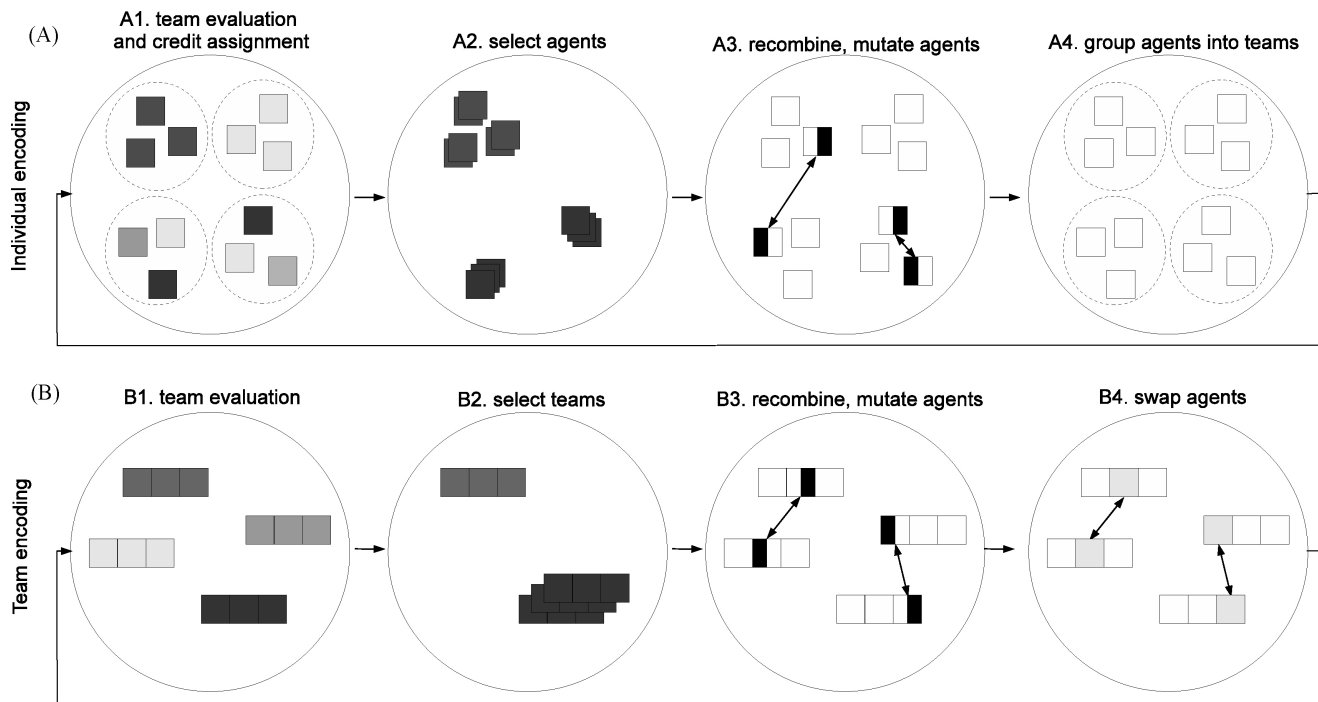


Fig. 1. Evolutionary algorithm cycle in the evolution of heterogeneous teams using (A) individual encoding and (B) team encoding. (A) Individual encoding: the evolutionary algorithm operates on a population of genotypes, each encoding one agent (depicted as squares). (A1) The agents are grouped into teams and their performance is evaluated together. Each agent must be assigned a fitness value (the credit assignment problem [12], [13], [90]). A straightforward way of addressing this problem is to distribute the team's performance score equally among the team members (top two teams in A1). This is known as global reward [12] or team-level selection [13]. Alternatively, the agents can be assigned with different fitness values, proportional to their personal contribution toward the team's performance score (bottom two teams in A1). This is known as local reward [12] or individual-level selection [13]. For simplicity, global and local rewards are illustrated together in one population, but typically only one kind of reward is used with all teams from the population. Next, (A2) the algorithm proceeds with the selection of the agents according to their fitness values. In (A1), a darker color denotes higher fitness values, which translates into more copies of fitter agents (A2). Afterward, in (A3) the agents' genotypes are recombined or mutated. Finally, in (A4) the agents are grouped into new teams before the performance evaluation and the algorithm starts over with (A1). (B) Team encoding: the evolutionary algorithm operates on a population of genotypes, each encoding one entire team (a square depicts a part of the team's genome corresponding to one agent). Consequently, there is no need for an extra step of grouping agents into teams and the algorithm begins at once with (B1) the evaluation of the teams' performance. There is no credit assignment problem, because team encoding implies that all agents in the team share the same fitness. Next, in (B2) the teams are selected according to their fitness values. In (B1), a darker color denotes higher fitness values, which translates into more copies of fitter teams (B2). Then, in (B3) the team's genotypes are subject to mutation or agent recombination. Finally, in (B4) entire agents may be swapped between the teams. This concludes the cycle and the algorithm starts over with (B1). Note that swapping agents between teams in team encoding conceptually corresponds to grouping of agents into teams in individual encoding.

De Jong and Spears [30], but not in the context of team evolution. In addition, Miconi [14] and Nakashima *et al.* [31] used a simple 1-point or 2-point crossover to recombine the teams' genotypes. This approach lends itself to the category of restricted crossovers. During the crossover, it usually swaps entire agents between teams, with the exception of the agents that happen to be placed on the crossover cutting points. The 1-point and 2-point crossovers, and generally  $n$ -point crossover, have been shown to have less exploratory power than uniform crossover [29].

A few attempts have been made to quantify the efficiency of agent recombination for the evolution of teams, but the studies are inconclusive [12]. Some authors advocate using restricted agent recombination [16], [17] and some are proponents of free agent recombination [20]. The problem of evolving team compositions has been addressed only for genetic programming by Hara and Nagao with automatically defined groups [32] and by Bongard with the Legion system [15]. In these two approaches, both the agents and their partitioning into separate subteams are evolved together. These approaches are promising, but designed for the tree representation typical for genetic programming. In addition, the authors considered only




Fig. 2. Applying four crossover operators on the genotype of two teams of four agents each. Black and white segments of teams' genotypes correspond to agents taking part in the crossover. Gray segments of teams' genotypes correspond to agents not taking part in the crossover. The crossover can either *recombine* the genetic material from two agents from parenting teams (top row) or *swap* entire agents between parenting teams (bottom row). In addition, the crossover may be *restricted* to act only on agents on corresponding positions in parenting teams (left column) or it may be *free* to act on any agents from parenting teams (right column).

a small number of evolving subteams (maximum number of subteams: six in [32] and three in [15]). Importantly, there is no study that quantifies the efficiency of agent swapping for the evolution of teams.

Operators similar to RAS and FAS have been proposed in different fields of evolutionary computation. Agent swapping

is equivalent to root crossover in genetic programming, which swaps whole trees instead of subtrees [23]. Also in genetic programming, Koza studied operators altering the architecture of a multipart program, which he used with automatic defined functions [33]. These operators included branch duplication and deletion, which, in the context of evolving team composition, would translate into agent deletion and duplication, respectively. In gene expression programming, Ferreira proposed gene recombination and gene transposition [34]. Gene recombination swaps entire genes between the parenting genotypes, and thus is equivalent to RAS. Gene transposition overwrites one gene with a copy of another gene, and thus is equivalent to agent deletion and duplication performed jointly. Finally, in the field of evolutionary strategies, Sebag and Schoenauer proposed mutation by imitation [35]. With imitation, the probability of an allele's mutation depends on whether the allele is the same or different from the corresponding alleles in some of the best or worst individuals in the population. Consequently, mutation by imitation might be considered a sort of restricted recombination.

### III. EXPERIMENTAL METHOD

#### A. RAS and FAS

We evolve teams of agents using team encoding and study RAS [Fig. 3(a)] and FAS [Fig. 3(b)]. We implement RAS as a uniform crossover [28], which exchanges the  $i$ th agent from the first team with the  $i$ th agent from the second team, with probability  $p$  for each agent. In an additional experiment, we tested RAS implemented as a 1-point and 2-point crossover (see Appendix A).

We implement FAS as a shuffle-uniform crossover [28], which exchanges the  $i$ th agent from the first team with the  $S(i)$ th agent from the second team, with probability  $p$  for each agent.  $S$  is a random permutation of integers from 1 to  $M$ , where  $M$  is the number of agents in the team. RAS may be considered a specific case of FAS where the sequence  $S$  is set to  $1, 2, \dots, M$ . For the sake of simplicity, we consider a situation with only two parenting teams, but the operators can scale to any number of parenting teams.

In addition, we theoretically investigate the connections between RAS, FAS, and other operators that alter the team composition inspired by architecture-changing operators [33] used in genetic programming and gene expression programming (see Appendix H).

#### B. Problem Formulation

We experimentally compare RAS and FAS in the evolution of agent teams facing the problem of finding the optimal team composition (i.e., the division of a team into groups of identical agents). Our aim is to mimic a situation when a team needs to display a certain composition of agents in order to achieve top performance [36]. For example, multiple robots may be more efficient if they compose distinct groups, each focusing on a different task [37]. In biology, this process is known as division of labor (see [38]–[40] for comprehensive reviews). For example, in many species of honey bees and




Fig. 3. RAS (left) and FAS (right) applied on two teams of eight agents. In the presented example, the probability of agent swapping was  $p = 0.5$ ; black lines denote the exchanged agents and the gray lines denote the agents that happen to be kept in their original team. The random permutation used to reshuffle the agents before the swap in (b) is  $(3, 1, 4, 8, 2, 5, 7, 6)$ .

ants, some workers forage for food, others care for the brood, and others perform maintenance work in the nest [41]–[44].

We consider a team consisting of  $M$  agents. There are  $D$  distinct types of agents. Thus, there are overall  $D^M$  different teams possible. We assume that the optimal team contains  $K$  distinct agents, each repeating  $R$  times (for simplicity, we assume  $M = K \cdot R$ ). Thus, the optimal composition of a team is defined as  $K$  groups of  $R$  agents, where the agents between the groups are different and the agents within a group are identical. The optimal team is homogeneous for  $K = 1$  and  $R = M$ , the optimal team is heterogeneous for  $K = M$  and  $R = 1$ , and the optimal team is hybrid [12] (also called partially heterogeneous [13]) for  $1 < K < M$  and  $1 < R < M$ .

For a real-life problem, the optimal team composition and, consequently, the values of  $K$  and  $R$  are unknown. They are discovered by means of artificial evolution, driven by a fitness function  $F$  that measures team performance for a given problem. It should be expected that team performance is correlated with the composition of the team. Thus, we set  $K$  and  $R$  *a priori* and we define a fitness function  $f$  that directly depends on the proportion of proper agents in each of the  $K$  groups  $f = \sum_{j=1}^K \min(\frac{R}{M}, x_j)$ . The value  $x_j$  is the fraction of agents from a team that belong to the  $j$ th group and operator  $\min(a, b)$  takes a value  $a$  if  $a < b$ , and  $b$  otherwise.

For example, consider three agents: A, B, and C. Let  $(x, y, z)$  denote the number of agents A, B, and C, respectively, in a team. The team size is set to six (i.e.,  $x + y + z = 6$ ), and the optimal team consists of three agents A and three agents B (i.e.,  $(3, 3, 0)$ ). Exactly one team  $(0, 0, 6)$  has the lowest performance  $f = 0$ . Exactly one team  $(3, 3, 0)$  has the maximal performance  $f = 1$ . For instance, teams  $(1, 1, 4)$ ,  $(0, 2, 4)$ , and  $(2, 0, 4)$  all have performance  $f = 0.33$ . Teams  $(6, 0, 0)$ ,  $(3, 0, 3)$ , and  $(1, 2, 3)$  all have performance  $f = 0.5$ . Teams  $(2, 2, 2)$ ,  $(1, 5, 0)$ , and  $(1, 3, 2)$  all have performance  $f = 0.67$ .

The fitness function  $f$ , although not directly applicable to real-life problems, lets us test the efficiency of RAS and FAS in the evolution of team composition under controlled conditions. The advantage of our approach is the ease of generating instances of various complexities and sizes. This enables us to perform systematic studies and thus draw statistically significant conclusions. We believe that the presented formulation provides a good abstraction of the core properties

of many problems, in which one needs to optimize the team composition in a multiagent system. We further elaborate on this issue in Section V. In addition, we show the applicability of our results by validating the efficiency of RAS and FAS in the evolution of decentralized controllers in a task-allocation problem (see Appendix G).

### C. Evolutionary Experiments

We evolve teams of agents in three treatments: 1) using RAS for all generations of the evolutionary algorithm; 2) using FAS for all generations of the evolutionary algorithm; and 3) using FAS for the first half and RAS for the second half, of all generations of the evolutionary algorithm. We also investigated alternative ways of combining FAS and RAS (see Appendix F).

We compare RAS and FAS under three conditions, where the optimal team is composed of 1000, 100, and 10 groups of 1, 10, and 100 identical agents, respectively, ( $K \in \{1000, 100, 10\}$  and  $R \in \{1, 10, 100\}$ ). Overall, this makes a total of nine experimental lines (3 treatments  $\times$  3 conditions). Each experimental line is replicated ten times. In all numerical experiments, we use populations of 1000 teams of  $M = 1000$  agents each. Population size and team sizes are kept constant across generations. The number of all distinct types of agents is set to  $D = 10\,000$ . Each evolutionary run lasts for 1000 generations, with the exception of an additional experiment, where the number of generations is set to 2000 (see Appendix G).

The software testbed has been implemented with the help of ECJ framework [45]. The numerical experiments have been run on the Pleiades cluster at École Polytechnique Fédérale de Lausanne, Lausanne.

### D. Genetic Architecture, Selection, and Reproduction

A team's genotype consists of 1000 alleles (one allele per agent), which are integers from 1 to 10 000. At the first generation of each evolutionary run, each of the  $1000 \times 1000$  alleles is independently set to a random integer value between 1 and 10 000 with uniform distribution. Teams are evaluated in the collaborative task (i.e., how similar the team's composition and the optimal composition are) and assigned a performance (see Section III-B). To construct the 1000 teams of the following generation, we select 500 times two teams. Each parent is independently selected from the current population using tournament selection with tournament size set to 2, with the exception of an additional experiment, where tournament size is set to 7 (see Appendix D). The two selected teams are reorganized with FAS or RAS, which results in two new teams that are added to the next generation's population. We use RAS and FAS with the probability of exchanging the agents between two teams set to  $p = 0.5$  for each position in a team, with the exception of an additional experiment, where  $p = 0.2$  is used (see Appendix B). Note that  $p = 0.5$  is the highest possible value, because swapping agents between teams A and B, with probability  $0.5 < p \leq 1$ , is equivalent to swapping agents between teams B and A with probability  $1 - p$ . The newly added teams are not subject to mutation, with the exception of an additional experiment, where each allele is independently




Fig. 4. Box and Whisker plots showing the mean performance of 1000 teams evolved with (a) RAS and (b) FAS (ten replicates). Teams consisted of 1000 agents, which needed to display three different compositions (1000, 100, and 10 groups of identical agents) in order for the team to achieve the optimal performance.

set to a random integer value between 1 and 10 000 (uniform distribution) with a probability 0.001 (see Appendix C). We do not use mutations in the main experiments for two reasons. First, our intention is to investigate the evolutionary dynamics of RAS and FAS. Thus, to get clearer results on the effects of agent swapping, we do not use the mutation. Second, with 10 000 different agents, the population of  $1000 \times 1000$  agents already contains each agent on average 100 times. Thus, the introduction of the innovative genetic material during evolution should not be necessary, if the agent swapping does efficiently compose optimal teams. Finally, one of the 1000 new teams of the following generation is randomly chosen, discarded, and replaced by an exact copy of the best team from the current generation (i.e., elitism of size 1).

### E. Statistical Analysis

To compare the teams evolved with RAS and FAS, we average, for each experimental line and replicate, team performance over 1000 teams (ten replicates) at generation 1000. We also report the best team performance from 1000 teams (ten replicates) (see Appendix E). We explain the differences in team performance in terms of variation between teams that RAS and FAS introduced into the population. To this aim, we calculated the standard deviation of teams' performance in a population for each of the 1000 generations. In particular, we compared the convergence time (i.e., the number of generations until the measured standard deviation reached zero) and the final variation level between the teams (i.e., the value of the measured standard deviation at generation 1000). To explain the differences in variation between teams, we compared the proportion of corresponding positions that in all teams contained only agents that could not become members of the optimal team, and we compared the proportion of agents in teams that could become members of the optimal team, averaged over 1000 teams. The last two characteristics were calculated every 25 generations, due to high demand for resources of these calculations (i.e., CPU and disk space). Statistical significance within multiple experimental lines was determined with the Kruskal–Wallis test (nonparametric one-way analysis of variance). Statistical significance between a pair of experimental lines was determined with the Wilcoxon test (rank sum test for equal medians).




Fig. 5. Mean variation [ $\pm$  standard deviation (s.d.) in gray] between the 1000 teams in a population (ten replicates), measured for all 1000 generations. The variation between the teams was quantified with the standard deviation of teams' performance in a population for (a) RAS and (b) FAS, in each of the three conditions (10, 100, and 1000 groups of identical agents in the optimal composition).

#### IV. RESULTS

We analyze the first two treatments (RAS and FAS) in Sections IV-A and IV-B and the third treatment (mixed FAS/RAS) in Section IV-C.

##### A. RAS and FAS

There were important differences in team performance between the three conditions (1000, 100, and 10 groups) at the 1000th generation for both treatments [RAS: Fig. 4(a), Kruskal–Wallis test,  $df = 2$ ,  $p < 0.001$ ; FAS: Fig. 4(b), Kruskal–Wallis test,  $df = 2$ ,  $p < 0.001$ ]. With RAS, the highest team performance was for 1000 groups (100 groups:  $-27.4\%$ ; 10 groups:  $-69.8\%$ ; three pairwise Wilcoxon tests  $df = 9$ ,  $p < 0.001$ ). In contrast, with FAS, the highest team performance was for ten groups (100 groups:  $-9\%$ ; 1000 groups:  $-35.4\%$ ; three pairwise Wilcoxon tests  $df = 9$ ,  $p < 0.001$ ).

The performance difference between conditions in treatments with RAS and FAS was caused by convergence issues, which were different for each of the two treatments. With RAS, the performance difference was associated with the convergence time, i.e., the number of generations until the standard deviation of team performance in the population reached zero [Fig. 5(a), mean  $\pm$ s.d. generations for 1000 groups:  $334.3 \pm 14.6$ ; 100 groups:  $294.2 \pm 36.7$ ; 10 groups:  $113.5 \pm 3.1$ ; Kruskal–Wallis test,  $df = 2$ ,  $p < 0.001$ ; three pairwise Wilcoxon tests  $df = 9$ ,  $p < 0.01$ ]. Therefore, the evolutionary algorithm using RAS suffered from premature convergence. In contrast, with FAS, the performance difference was associated with the variation between the evolved teams, i.e., the value of the standard deviation of team performance in the population at generation 1000 [Fig. 5(b); Kruskal–Wallis test,  $df = 2$ ,  $p < 0.001$ ; three pairwise Wilcoxon tests  $df = 9$ ,  $p < 0.001$ ]. Therefore, the evolutionary algorithm using FAS suffered from insufficient convergence.

With RAS, premature convergence was detrimental to team performance because it leads to the disappearance of agents required in the optimal team on corresponding positions in teams across the entire population [Fig. 6(a), top row]. This was not the case for treatments with FAS [Fig. 6(b), top row]. Consequently, in treatments with RAS, the teams contained




Fig. 6. (Top) Mean proportion ( $\pm$ s.d. in gray) of corresponding positions that contained in all 1000 teams only the agents that could not become members of the optimal team. (Bottom) Mean proportion ( $\pm$ s.d. in gray) of agents in a team that could become members of the optimal team, averaged over 1000 teams. The populations were analyzed every 25 generations over ten replicates for (a) RAS and (b) FAS, in each of the three conditions (10, 100, and 1000 groups of identical agents in the optimal composition).

agents that could not become members of the optimal team [Fig. 6(a), bottom row]. With FAS, insufficient convergence was detrimental to team performance because it prevented the formation of the optimal composition. This is supported by the fact that team performance stayed low, even though the teams contained only the agents required in the optimal team [Fig. 6(b), bottom row]. We illustrate this process with a simple thought experiment. Consider a population consisting of two identical teams each containing agents A and B (in that order). In such conditions, RAS may swap agent A with agent A only, and agent B with agent B only. This does not change team compositions in the population and thus does not affect the teams' performance. In contrast, FAS may at some point swap agent A with B, which would result in one team containing both agents A and the second team containing both agents B. Consequently, FAS may destroy favorable team compositions, even when entire population contains only optimal teams.

We performed a sensitivity analysis to see how our results were affected by lower probability of agent swapping, by mutations, and by higher selection pressure. Lower probability of agent swapping had a small detrimental effect on team performance in both treatments and was more marked for RAS (see Appendix B). In treatments with RAS, using the mutation counterbalanced, to some extent, premature convergence [compare Figs. 5 and 17], which led to an increase in team performance. In contrast, with FAS, the mutation introduced additional undesirable variation between the teams and thus had a weak detrimental effect on the performance (see Appendix C). In treatments with RAS, stronger selection increased the detrimental effect of premature convergence on team performance. With FAS, stronger selection did not overcome the problem of insufficient convergence and had no effect on the performance of the evolved teams (see Appendix D).




Fig. 7. Box and Whisker plots showing the mean performance of 1000 teams evolved with (a) RAS and (b) FAS (ten replicates). Teams consisted of 1000 agents, which needed to display three different compositions (10, 100, and 1000 groups of identical agents) in order for the team to achieve the optimal performance. The populations were initialized randomly with a bias. For each of the three conditions and for both treatments, the initial population always contained on average 10% of agents that could become members of the optimal team.

We tested alternative implementations of RAS, i.e., 1-point and 2-point crossover. The results indicate that uniform crossover is better suited for RAS than  $n$ -point crossover (see Appendix A). Here, we presented the mean team performance in the population (see Figs. 4 and 7). We also report the best team performance in the population (averaged over ten replicates), which shows not to be qualitatively different from the mean team performance (see Appendix E).

### B. RAS and FAS: Disparities in the Initial Population

In the previous experiments, the agents in the teams were initialized randomly with uniform distribution (i.e., each initialized agent was set to be one of the 10000 possible agents with probability 0.0001). Consequently, the proportion of agents in the initial population that could become members of the optimal team varied between conditions. In the first condition with 1000 groups, the initial population contained on average 10% of agents that could become members of the optimal team. In the second condition with 100 groups, the initial population contained 1% of such agents. In the third condition with 10 groups, the initial population contained only 0.1% of such agents. We tested if these disparities in the initial population influenced the difference in team performance between conditions for both treatments. To this aim, we performed a control experiment in which the populations were initialized randomly with a bias. For each of the three conditions, the initial population always contained on average 10% of agents that could become members of the optimal team.

In treatments with RAS and without disparities in the initial population, there was an important increase in team performance for conditions with 10 and 100 groups in the optimal composition [compare Figs. 4(a) and 7(a); two Wilcoxon tests,  $df = 9$ ,  $p < 0.001$ ]. By contrast, there was no significant change over the 1000 generations for 1000 groups [compare Figs. 4(a) and 7(a); Wilcoxon test  $df = 9$ ,  $p = 0.47$ ]. This suggests that with RAS team, performance primarily depended on the number of agents that could become members of the optimal team. The performance of the evolved teams was higher with more such agents in the initial population.




Fig. 8. Mean performance ( $\pm$ s.d. in gray) of 1000 teams evolved with (a) RAS, (b) FAS, and (c) mixed FAS/RAS (10 replicates). In (c), the teams were evolved with FAS in generations from 1 to 500 (inclusive) and with RAS in generations from 501 to 1000. The black solid vertical line at generation 500 marks the transition from using FAS to using RAS. Teams consisted of 1000 agents, which needed to display three different compositions (10, 100, and 1000 groups of identical agents) in order for the team to achieve the optimal performance.

In treatments with FAS, in contrast to RAS, there were no important differences in performance of teams evolved with and without disparities in the initial population for each of the three conditions [compare Figs. 4(b) and 7(b), three pairwise Wilcoxon tests  $df = 9$ ,  $p > 0.3$ ]. This suggests that with FAS team performance primarily depended on the optimal composition. The performance of the evolved teams was higher with a lower number of groups in the optimal team.

### C. Combining FAS and RAS

In the treatments with RAS and FAS, the exploration and exploitation aspects of the evolutionary algorithm were not well balanced. With RAS, the evolutionary algorithm suffered from premature convergence, whereas with FAS, it suffered from insufficient convergence. Consequently, in both cases, only suboptimal solutions have evolved [Fig. 8(a) and (b)]. In order to overcome the limitations of both RAS and FAS, we combined the two swapping methods in a complementary way: for the first 500 generations, FAS was used and for the next 500, generations RAS was used. We expected this would allow us to efficiently explore the search space first and then to exploit it. We tested the efficiency of the mixed approach in the evolution of team composition for the three conditions (1000, 100, and 10 groups in the optimal composition).

There were important differences in team performance between the treatments (RAS, FAS, FAS/RAS) at the 1000th generation for each of the three conditions [Fig. 8, three Kruskal–Wallis tests,  $df = 2$ ,  $p < 0.001$ ]. For each of the three conditions, the performance was higher with FAS/RAS than with RAS alone [Fig. 8(a) and (c), three Wilcoxon tests  $df = 9$ ,  $p < 0.001$ ] and with FAS alone [Fig. 8(b) and (c), three Wilcoxon tests  $df = 9$ ,  $p < 0.001$ ]. This was because in treatments with the mixed FAS/RAS approach, using FAS for the first 500 generations resulted in an abundance of agents that could become members of the optimal team [Fig. 6(b), bottom row]. These agents could then be efficiently used to compose optimal teams using RAS in the following 500 generations. In addition, for the first 500 generations, there was high variation between the evolving teams typical for FAS, which was then decreased by applying RAS [Fig. 9]. Overall,




Fig. 9. Mean variation ( $\pm$ s.d. in gray) between the 1000 teams in a population, quantified with the standard deviation of teams’ performance in a population. The teams were evolved in three different conditions (10, 100, and 1000 groups of identical agents in the optimal composition) with FAS in generations from 1 to 500 (inclusive) and with RAS in generations from 501 to 1000 (ten replicates). The black solid vertical line at generation 500 marks the transition from using FAS to using RAS.

this mixed approach overcame the convergence issues of both FAS and RAS.

We fixed the moment of switch from FAS to RAS at the middle of the evolutionary process (500th generation). But, if needed, one may envision a dynamic switch between the agent swapping methods, which should be performed as soon as the variation between teams stabilizes (here around the 250th generation) [Fig. 9]. In addition, we tested alternative implementations of combining RAS and FAS, by applying both operators simultaneously but with different and varying probabilities of agent swapping. The results indicate that RAS and FAS should be used exclusively in order to balance the exploration and exploitation aspects of the evolutionary algorithm (see Appendix F).

## V. DISCUSSION

We considered a situation when an optimal team consisted of equally sized groups of identical agents, but the results allow us to speculate how the RAS and FAS compare when groups differ in sizes. In treatments with RAS, team performance depended on the contents of the initial population and not on the actual optimal composition. Consequently, limitations and analysis presented in this paper for RAS should also hold for optimal team composition with groups of different sizes. This is because no link should be expected between the number of agents of a specific type in the initial population and in the optimal composition.

In contrast, with FAS, team performance depended on the optimal team composition. FAS introduced a variation between team compositions, which was detrimental to team performance. The level of this variation depended on the condition. For example, the variation was the highest and hence team performance was the lowest, when the optimal team was composed of many small groups (here 1000 groups of one agent). This was because, with FAS, it is more probable for a team to deviate from the optimum when the optimal composition consists of small groups than when it consists of big groups. We illustrate this property with a simple thought experiment. Consider two teams of four agents, both displaying the same optimal composition that consists of four

groups of one agent. There are 16 possible free agent swaps between the two teams. Only four of them, the ones that happen on the corresponding positions, maintain the optimal composition in the teams. In contrast, consider two teams of four agents, both displaying the same optimal composition that consists of two groups of two agents. Again, there are 16 possible free agent swaps between the two teams. But now there are eight swaps that maintain the optimal composition in the teams. Therefore, the performance of teams evolved with FAS depends on the size of groups in the optimal composition. But it does not depend on if these groups are of equal size or not. Consequently, limitations and analysis presented in this paper for FAS should also hold for optimal team composition with groups of different sizes.

Several general guidelines on applying RAS and FAS to real-life problems can be drawn. First, we compared RAS and FAS on a problem of evolving team compositions, whose formulation was highly general. The part of a team’s genotype that corresponded to a single agent was rather simplistic, i.e., a vector of one (main experiments) to a few numbers (see additional experiments in Appendix G). Consequently, our results are directly applicable whenever the goal is to optimize just a few control parameters per agent and not to create entire controllers for each agent from scratch. This is the case for many real-life applications, when the problem is often decomposed into two parts: the one solved manually and the one solved automatically. For example, in behavioral robotics, agents perform behaviors that have been implemented manually [46]–[48]. In bio-inspired scheduling, agents allocate themselves to tasks which they handle by calling existing routines [49]–[52]. In real-time strategy games, units perform pre-defined actions or entire sequences of actions [53], [54]. In ant colony optimization, the agents traverse the edges of a graph according to established algorithms [2], [55]. In all these cases, the agents choose behaviors, tasks, actions, and edges based on the values of a few control parameters that can be optimized automatically [56], [57]. In conclusion, our problem’s formulation, although general, conforms very well to practical applications, as diverse as robotics, task allocation, video games, and hyperheuristics (i.e., searching in the space of heuristics).

Three implementation issues related to RAS and FAS should be discussed. First, we considered the agent swapping operators to work on teams having linear genotypes (in contrast to, e.g., a tree-based representation popular in genetic programming). Both RAS and FAS can be easily adapted to other teams’ representations. In principle, FAS requires teams to be multisets of agents, i.e., sets in which elements may repeat. Whereas RAS requires teams to be sequences of agents. This is not much of a limitation because one can always create a sequence by imposing an ordering on a multiset. Note, however, that the ordering must remain constant during the evolution or, at least, it must change in the same way for all teams in the population. Otherwise, the agents could change their positions due to variations in the ordering. Consequently, there would be no restrictions on agents keeping their positions in teams and using RAS would make little sense.

Second, we evolved teams of constant size and thus we used genotypes of constant length. Both RAS and FAS could



be adapted to variable-length genotypes by restraining the operators to act only on common parts of the genotypes. Alternatively, one could use agent deletion and duplication that alters the team compositions similar to FAS (see Appendix H).

Third, for the mixed FAS/RAS operator, we also considered alternative implementations. We applied both operators simultaneously but with different and varying probabilities of agent swapping. The results discourage such an approach and suggest that RAS and FAS should be used exclusively. This is due to the disruptive character of FAS, which is strongly marked even for low probabilities of agent swapping (see Appendix F).

We focused on the evolution of team compositions using team encoding, i.e., when a single genotype encodes all individuals from one team. Team encoding implicitly assumes team level of selection (i.e., global reward, see Fig. 1). In contrast to individual level of selection, team level of selection omits pathologies caused by competition between team members (e.g., evolution of cheaters [58]–[60]), which could decrease the overall team performance. Therefore, team level of selection and, consequently, team encoding are advocated when the goal is to optimize the overall team’s performance [13]. Studying interagent interactions usually requires a game-theoretic perspective and is an interesting subject in itself (e.g., [61]), but was not in the scope of this paper.

Nevertheless, our results are also relevant for evolutionary algorithms using individual encoding (e.g., [13], [18], [24]–[27]). This is because the converging character of RAS and the disruptive character of FAS are general properties of these operators. Thus, premature and insufficient convergence would be to some extent marked with RAS and FAS, respectively, regardless of the encoding and the level of selection used. For example, with individual encoding, in order to assess performance, the agents are grouped into teams, often at random (e.g., [26], [62], and [13]). Random grouping in individual encoding introduces a constant variation between compositions of evolving teams, similar to FAS in team encoding. Consequently, random grouping shares the limitations of FAS and may hamper the evolutionary process whenever high level of genetic specialization between agents is required.

It should be noted that the evidence from our paper has three limitations, which translate into three directions of future research.

First, we assumed that a single change in team composition perfectly translates into a corresponding change in team performance. It might be desirable to consider epistatic and noisy fitness functions (e.g., [63]). We expect FAS to drive the evolution toward optimum even then because of the highly explorative nature of this method. We supported this claim by validating agent-swapping operators in a stochastic and dynamic problem of decentralized task allocation (see Appendix G). In this practical application, we showed that the teams evolved with FAS had higher performance than the teams evolved with RAS, probably due to low level of agent specialization that was required to solve the problem. This result is consistent with our previous analyses and shows that the conclusions are valid also for more complex problems.

Second, we assumed a situation with no locally optimal team compositions. In practice, the globally optimal team composition might be hidden in a part of the landscape with low average payoff (i.e., deceptive fitness landscape [64]). Such a fitness landscape could hamper the evolution of optimal teams, similar to other problems being solved with evolutionary algorithms [65], [66]. Fortunately, evolutionary algorithms frequently turn out to be excellent heuristics for most deceptive fitness landscapes [64].

Third, we focused on teams consisting of agents that did not evolve themselves. This allowed us to decouple the effects of agent recombination and agent swapping and to directly study the evolution of team compositions. We also validated RAS and FAS including mutations. This did not affect the conclusions drawn (see Appendixes C and G). Note that developing the methods that efficiently evolve large teams of complex genotypes for multiagent systems remains a big challenge. So far, other authors have focused on the evolution of complex controllers for small teams of agents (e.g., GP trees [16], [21] and neural networks [13], [26]) and have overlooked the evolution of team compositions (with the exception of [32] and [15], which consider only simple teams of up to six groups). Consequently, our investigation of evolving team compositions fills the gap and hopefully will help to address other challenging questions on optimizing multiple agents.

## VI. CONCLUSION

In this paper, we were able to quantify for the first time the efficiency of RAS and FAS in the evolution of team compositions using team encoding. The analysis and comparison between RAS and FAS revealed the limitations of both approaches. Our paper highlights and explains the convergence issues, which were detrimental to team performance. Using RAS resulted in an efficient evolution of team compositions, only if the population contained enough agents that could become members of the optimal team (but regardless of the optimal composition). In contrast, using FAS resulted in an efficient evolution of team compositions, only if optimal composition imposed a low requirement on genetic specialization of the agents (but regardless of the contents of initial population). In order to overcome these limitations, the results strongly suggested using a combination of the two methods of agent swapping, which were so far only used in isolation. The mixed approach balanced the exploration and exploitation aspects of the evolutionary algorithm and experimentally proved to be a superior strategy in the evolution of team compositions.

## APPENDIX A

### 1-POINT AND 2-POINT CROSSOVER

Miconi [14] and Nakashima *et al.* [31] used 1-point and 2-point crossover to recombine the teams’ genotypes. The  $n$ -point crossover might be considered a competitive way to implement RAS. It has been shown that  $n$ -point crossover has lower exploratory power than uniform crossover [29], [30]. Also, teams are sets of agents (formally multisets) and not sequences of agents, thus the disruptive nature of the




Fig. 10. Box and Whisker plots showing the mean performance of 1000 teams evolved with (a) 1-point crossover and (b) 2-point crossover (ten replicates). Teams consisted of 1000 agents, which needed to display three different compositions (10, 100, and 1000 groups of identical agents) in order for the team to achieve the optimal performance.

uniform crossover (i.e., breaking the schemata) [30] is not necessarily an undesirable feature. These two facts suggest that uniform crossover [21], [28] is a better choice than  $n$ -point crossover in the implementation of RAS. We found support for this claim, and we showed that in treatments with RAS implemented as 1-point and 2-point crossover team performance was lower [Fig. 10] than in treatments with RAS implemented as uniform crossover. This performance drop was associated with faster convergence [Fig. 11] with 1-point and 2-point than uniform crossover. Consequently, the proportion of corresponding positions that in all teams contained only agents that could not become members of the optimal team was higher [Fig. 12, top row] with 1-point and 2-point than with uniform crossover. The proportion of agents in teams that could become members of the optimal team was lower [Fig. 12, bottom row] with 1-point and 2-point than with uniform crossover.

In treatments with 1-point and 2-point crossovers, to construct the 1000 teams of the following generation, we selected, 500 times, two teams (tournament size was set to 2). Then, instead of reorganizing the two teams with uniform crossover, the teams' genotypes were crossed over with 1-point or 2-point crossovers. With 1-point crossover, one random locus  $1 \leq L \leq 1000$  was chosen and the parenting teams exchanged all agents on positions from 1 to  $L$ . With 2-point crossover, two random loci  $1 \leq L_1 \leq L_2 \leq 1000$  were chosen and the parenting teams exchanged all agents on positions from  $L_1$  to  $L_2$ . Other settings were the same as in Section III.

## APPENDIX B

### EFFECTS OF AGENT SWAPPING WITH PROBABILITY $p = 0.2$

To test the sensitivity of the results to different probability of swapping the agents between the parenting teams, we performed an additional experiment with this probability set to 0.2. Other settings were the same as in Section III. We found that the control experiment is in reasonably good agreement with the previously reported results with respect to team performance [Fig. 13], the variation between teams [Fig. 14], the proportion of corresponding positions that in all teams contained only agents that could not become members of the optimal team [Fig. 15, top row], and the proportion of agents




Fig. 11. Mean variation ( $\pm$ s.d. in gray) between the 1000 teams in a population (ten replicates), measured for all 1000 generations. The variation between the teams was quantified with the standard deviation of teams' performance in a population for (a) 1-point crossover and (b) 2-point crossover, in each of the three conditions (10, 100, and 1000 groups of identical agents in the optimal composition).




Fig. 12. (Top) Mean proportion ( $\pm$ s.d. in gray) of corresponding positions that contained, in all 1000 teams, only the agents that could not become members of the optimal team. (Bottom) Mean proportion ( $\pm$ s.d. in gray) of agents in a team that could become members of the optimal team, averaged over 1000 teams. The populations were analyzed every 25 generations over ten replicates for (a) 1-point crossover and (b) 2-point crossover, in each of the three conditions (10, 100, and 1000 groups of identical agents in the optimal composition).

in teams that could become members of the optimal team [Fig. 15, bottom row].

## APPENDIX C

### EFFECTS OF MUTATION

With RAS, the evolutionary algorithm suffered from premature convergence. Therefore, using mutation with RAS should have a positive effect on the exploration aspect of the evolutionary algorithm and, thus, translate into an increase in team performance. In contrast, with FAS, the evolutionary algorithm suffered from insufficient convergence. Consequently, using mutation with FAS should have a negative effect on the exploitation aspect of the evolutionary algorithm and, thus, translate into a decrease in team performance. We found support for both these claims in an additional experiment where the teams' genotypes were subject to mutation, i.e.,

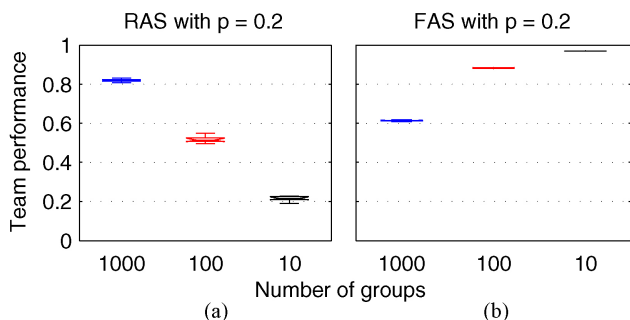


Fig. 13. Box and Whisker plots showing the mean performance of 1000 teams evolved with (a) RAS and (b) FAS (ten replicates). Teams consisted of 1000 agents, which needed to display three different compositions (10, 100, and 1000 groups of identical agents) in order for the team to achieve the optimal performance. For both treatments (RAS and FAS), the agents were swapped with probability 0.2.




Fig. 14. Mean variation ( $\pm$ s.d. in gray) between the 1000 teams in a population (ten replicates), measured for all 1000 generations. The variation between the teams was quantified with the standard deviation of teams' performance in a population for (a) RAS and (b) FAS, in each of the three conditions (10, 100, and 1000 groups of identical agents in the optimal composition). For both treatments (RAS and FAS), the agents were swapped with probability 0.2.

each allele of the newly added teams to the population was randomly set to a value between 1 and 10000 with a probability 0.001. Other settings were the same as in Section III.

In treatments with RAS, using the mutation counterbalanced, to some extent, premature convergence [Fig. 17(a)] and its negative effect on team compositions, with respect to the proportion of corresponding positions that contained, in all teams, only agents that could not become members of the optimal team [Fig. 18(a), top row] and to the proportion of agents in teams that could become members of the optimal team [Fig. 18(a), bottom row]. This translated into a higher team performance [Fig. 16(a)] than in treatments with RAS without mutation. In contrast, in treatments with FAS, the mutation increased the variation between the teams [Fig. 17(b)], which had a detrimental effect on evolving team compositions [Fig. 18(b), top row and Fig. 18(b), bottom row] and, consequently, on team performance [Fig. 16(b)].

#### APPENDIX D EFFECTS OF SELECTION PRESSURE

We performed a sensitivity analysis and investigated how our results were affected by increasing the selection pressure. To this aim, we used a tournament size set to 7. Other

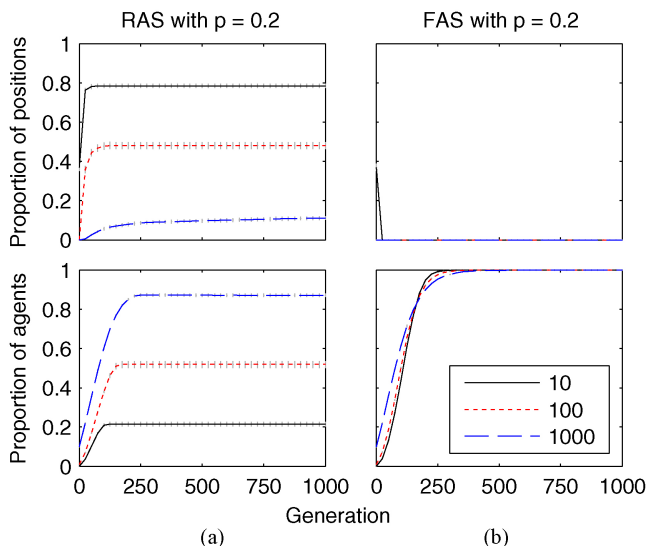


Fig. 15. (Top) Mean proportion ( $\pm$ s.d. in gray) of corresponding positions that contained in all 1000 teams only the agents that could not become members of the optimal team. (Bottom) Mean proportion ( $\pm$ s.d. in gray) of agents in a team that could become members of the optimal team, averaged over 1000 teams. The populations were analyzed every 25 generations over ten replicates for (a) RAS and (b) FAS, in each of the three conditions (10, 100, and 1000 groups of identical agents in the optimal composition). For both treatments (RAS and FAS), the agents were swapped with probability 0.2.

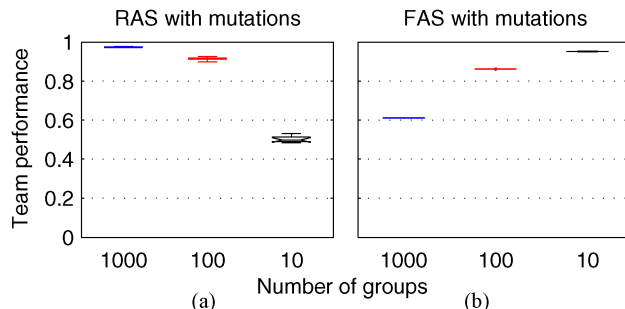


Fig. 16. Box and Whisker plots showing the mean performance of 1000 teams evolved with (a) RAS and (b) FAS (ten replicates). Teams consisted of 1000 agents, which needed to display three different compositions (10, 100, and 1000 groups of identical agents) in order for the team to achieve the optimal performance. For both treatments (RAS and FAS), each allele in a team's genotype was mutated with probability 0.001.

settings were the same as in Section III. We found that with RAS, stronger selection increased premature convergence [Fig. 20(a)], which translated into lower team performance [Fig. 19(a)], than in treatments with tournament size set to 2 (see Section IV). This was because the increased premature convergence led to higher proportion of corresponding positions that in all teams contained only agents that could not become members of the optimal team [Fig. 21(a), top row] and to a lower proportion of agents in teams that could become members of the optimal team [Fig. 21(a), bottom row].

In contrast, with FAS, stronger selection had no effect on the performance of the evolved teams [Fig. 19(b)] because it did not overcome the problem of insufficient convergence [Fig. 20(b)]. Consequently, the results for treatments with FAS

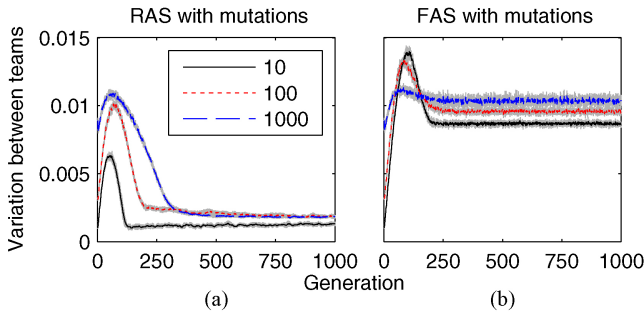


Fig. 17. Mean variation ( $\pm$ s.d. in gray) between the 1000 teams in a population (ten replicates), measured for all 1000 generations. The variation between the teams was quantified with the standard deviation of teams' performance in a population for (a) RAS and (b) FAS, in each of the three conditions (10, 100, and 1000 groups of identical agents in the optimal composition). For both treatments (RAS and FAS), each allele in a team's genotype was mutated with probability 0.001.

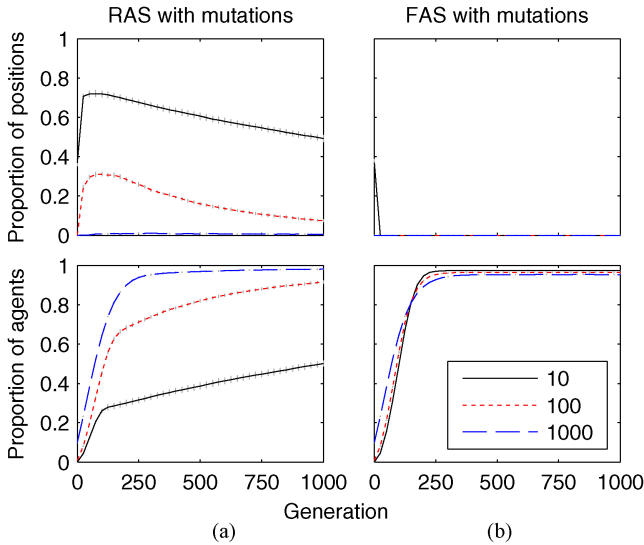


Fig. 18. (Top) Mean proportion ( $\pm$ s.d. in gray) of corresponding positions that contained in all 1000 teams only the agents that could not become members of the optimal team. (Bottom) Mean proportion ( $\pm$ s.d. in gray) of agents in a team that could become members of the optimal team, averaged over 1000 teams. The populations were analyzed every 25 generations over ten replicates for (a) RAS and (b) FAS, in each of the three conditions (10, 100, and 1000 groups of identical agents in the optimal composition). For both treatments (RAS and FAS), each allele in a team's genotype was mutated with probability 0.001.




Fig. 19. Box and Whisker plots showing the mean performance of 1000 teams evolved with (a) RAS and (b) FAS (ten replicates). Teams consisted of 1000 agents, which needed to display three different compositions (10, 100, and 1000 groups of identical agents) in order for the team to achieve the optimal performance. For both treatments (RAS and FAS), the tournament size in the tournament selection was set to 7.

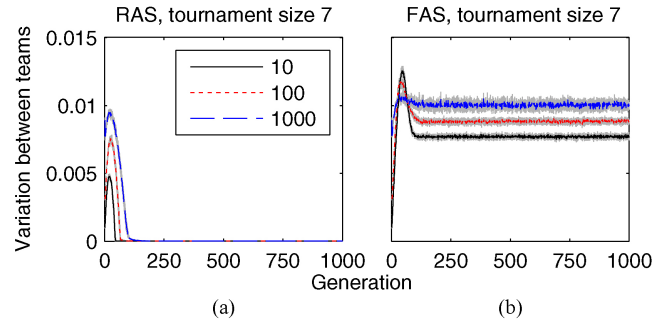


Fig. 20. Mean variation ( $\pm$ s.d. in gray) between the 1000 teams in a population (ten replicates), measured for all 1000 generations. The variation between the teams was quantified with the standard deviation of teams' performance in a population for (a) RAS and (b) FAS, in each of the three conditions (10, 100, and 1000 groups of identical agents in the optimal composition). For both treatments (RAS and FAS), the tournament size in the tournament selection was set to 7.

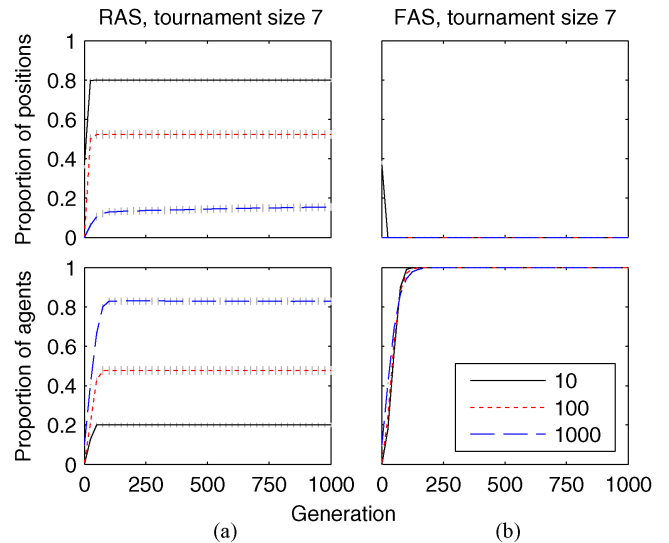


Fig. 21. (Top) Mean proportion ( $\pm$ s.d. in gray) of corresponding positions that contained in all 1000 teams only the agents that could not become members of the optimal team. (Bottom) Mean proportion ( $\pm$ s.d. in gray) of agents in a team that could become members of the optimal team, averaged over 1000 teams. The populations were analyzed every 25 generations over ten replicates for (a) RAS and (b) FAS, in each of the three conditions (10, 100, and 1000 groups of identical agents in the optimal composition). For both treatments, RAS and FAS, the tournament size in the tournament selection was set to 7.

with the two strengths of selection pressure are in agreement with respect to the proportion of corresponding positions that in all teams contained agents that could not become members of the optimal team [Fig. 21(b), top row] and the proportion of agents in a team that could become members of the optimal team [Fig. 21(b), bottom row].

### APPENDIX E

#### BEST TEAM PERFORMANCE IN THE POPULATION

We report here the performance of the best team at generation 1000th over ten replicates [Fig. 22] and compare it with mean team performance. To this aim, we calculated over ten replicates the mean  $\pm$ s.d. ratio  $r$  between the value of mean performance in the population and the value of the best

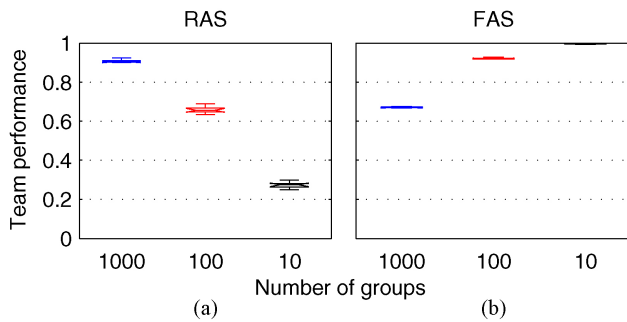


Fig. 22. Box and Whisker plots showing the best performance from 1000 teams evolved with (a) RAS and (b) FAS (ten replicates). Teams consisted of 1000 agents, which needed to display three different compositions (10, 100, and 1000 groups of identical agents) in order for the team to achieve the optimal performance.

performance. There were no differences between mean and best team performance for treatments with RAS ( $r = 1 \pm 0$  for all three conditions). There were small differences between mean and best team performance for treatments with FAS ( $r$  for 1000 groups:  $0.931 \pm 0.004$ ; 100 groups:  $0.955 \pm 0.003$ ; 10 groups:  $0.971 \pm 0.001$ ). These differences between mean and best team performance correspond well to the variation between teams at the 1000th generation (0 for RAS, and approximately 0.01 for FAS, see Fig. 5).

## APPENDIX F

### ALTERNATIVE METHODS OF COMBINING RAS AND FAS

Earlier, we combined FAS with RAS by applying them from generation 1 to 500 and from generation 501 to 1000, respectively. Here, we investigated two alternative methods that apply both operators together, but with different probabilities of agent swapping [Fig. 23, top row]. With the first method, the probability of agent swapping with FAS ( $p_{FAS}$ ) was set to  $= 0.5$  at generation 1 and decreased with a constant rate to 0 at generation 1000. With the second method, the  $p_{FAS}$  was set to 0.5 from generation 1 to 250, then it decreased with a constant rate to 0 until generation 750 and was fixed to 0 until generation 1000. In both methods, the probability of agent swapping with RAS ( $p_{RAS}$ ) was set to  $0.5 - p_{FAS}$ . Other settings were the same as in Section III.

The results for both alternative methods were discouraging [Fig. 23, middle row] because the disruptive character of FAS was strongly marked even for low probabilities of agent swapping. Consequently, the first alternative method did not converge [Fig. 23(a), bottom row] and the second alternative method started converging only after  $p_{FAS}$  decreased to 0 [Fig. 23(b), bottom row]. Therefore, in order to balance well the exploration and exploitation aspects in the evolution of team compositions, one should use FAS and RAS exclusively.

## APPENDIX G

### APPLICATION OF RAS AND FAS TO THE EVOLUTION OF DECENTRALIZED CONTROLLERS FOR TASK ALLOCATION

*Methods:* We validated the agent-swapping operators (RAS and FAS) in a complex problem of evolving decentralized




Fig. 23. (Top row) The probabilities of agent swapping with two alternative methods of combining (a) RAS and (b) FAS. (Middle row) Mean performance ( $\pm$ s.d. in gray) of 1000 teams. Teams consisted of 1000 agents, which needed to display three different compositions (10, 100, and 1000 groups of identical agents) in order for the team to achieve the optimal performance. (Bottom row) Mean variation ( $\pm$ s.d. in gray) between the 1000 teams. The variation between the teams was quantified with the standard deviation of teams' performance in a population. Both the mean performance and the mean variation were measured for all 1000 generations (ten replicates).

controllers for task allocation. We considered a team composed of 1000 agents that allocated themselves to two distinct tasks using thresholds-based models [49]–[52]. Agents receive information of the team's needs via commonly perceived stimuli. Then, agents respond to team's needs based on the values of their internal thresholds. Here, we optimized two different versions of the response thresholds models. First, with deterministic response threshold model (DTM) [67], [68], every agent had two thresholds corresponding to each of the two tasks. An agent performed the task with the highest positive difference between the stimulus and its own corresponding response threshold or remained idle if both of its thresholds were higher than the stimuli. If the difference between the stimulus and the agent's corresponding response thresholds was the same for all tasks, one of them was randomly chosen and performed by the agent. Second, with the extended response threshold model (ETM) [57], every agent had two thresholds corresponding to each of the two tasks and two weights corresponding to each of the two stimuli. An agent performed the task with the highest positive difference between the weighted stimulus and its own corresponding response threshold or remained idle if both of its thresholds were higher than the weighted stimuli. If the difference between the weighted stimulus and the agent's corresponding response threshold was the same for all tasks, one of them was randomly chosen and performed by the agent.

To quantify the teams' efficiency in task allocation, we used a stochastic agent-based simulation to model a situation in which agents had to perform two distinct tasks [57], [69]. Our

aim was to mimic situations with two vital tasks such as foraging and regulation. A team consisted of 1000 agents placed in an environment with an infinite number of two types of items: foraging and regulatory. The team’s lifespan was divided into 100 time-steps. At the beginning of each time-step, an agent was presented with two task stimuli, one for the foraging items and the other for the regulatory items. If there were no items in the base, the corresponding stimulus was set to its maximal intensity, which was 1. Otherwise, the intensity of the stimulus for each task was inversely proportional to the number of corresponding items in the base. The foraging stimulus at time-step  $t$  was equal to  $1 - 10^{-4} \cdot a^F(t-1)$ , where  $a^F(t-1)$  is the number of foraging items accumulated in the base at time step  $t-1$ . The regulatory stimulus at time-step  $t$  was equal to  $1 - 5 \cdot 10^{-3} \cdot a^R(t-1)$ , where  $a^R(t-1)$  is the number of regulatory items accumulated in the base at time step  $t-1$ . At each time step, every agent performed the chosen task (or stayed idle) according to the task allocation mechanism (DTM, ETM) considered in the experiment. At each time step, an agent had a probability of 0.1 to successfully collect one item corresponding to the task performed. At each time-step, the number of foraged items in the base were depleted by ten items if they were not depleted in the previous time-step, otherwise they were depleted by ten items with the probability of  $\frac{1}{2}$ . The same procedure was applied to independently deplete the regulatory items.

The team performance directly depended on the number of collected foraging items, but these were counted only when the number of regulatory items in the base was within predefined bounds (140–160 items). At the first time-step of a simulation, there were no items of the foraging and regulatory tasks in the base. The team performance  $f$  was calculated by adding the partial performance obtained at each time-step, with  $f = \sum_{t=1}^{100} f(t)$ , where the team performance at each time-step ( $f(t)$ ) was quantified as the number of items of the foraging task collected when the number of items of the regulatory task present in the base was between 140 and 160:  $f(t) = b(t) \cdot g^F(t)$ , where  $b(t) = 1$  if  $140 \leq g^R(i) \leq 160$  and  $b(t) = 0$  otherwise,  $g^F(t)$  represents the number of items foraged at time-step  $t$  and  $g^R(t)$  the number of items being regulated within the base at time-step  $t$ . Thus, if teams performed well in only one of the two tasks, their fitness was low. We normalized the resulting fitness values by 10 000, which is the expected amount of foraging items collected if all 1000 agents were foraging for all 100 time-steps with the probability of success equal to 0.1.

In experiments with DTM and ETM, each team had a genome consisting of 1000 parts, which corresponded to the agents. With DTM, each of these 1000 genome’s parts consisted of two thresholds, both ranging from  $-1$  to  $+1$  (8-bit encoding, 256 possible values with a resolution of  $\frac{1}{128}$ ). With ETM, each of the 1000 genome’s parts consisted of two thresholds and two weights, all ranging from  $-1$  to  $+1$  (8-bit encoding, 256 possible values with a resolution of  $\frac{1}{128}$ ).

We performed 2000 generations of artificial selection in 30 independent replicates for each of the two models (DTM and ETM), for two treatments (RAS and FAS), and in two conditions (without and with mutations). Overall, there were  $2 \times 2 \times 2 = 8$  experimental lines. In the experiments with

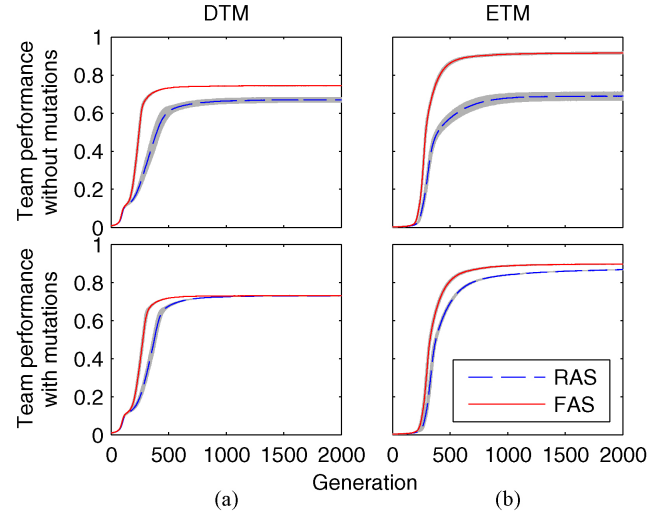


Fig. 24. Mean performance ( $\pm$ s.d. in gray) of 1000 teams evolved with RAS and FAS for two treatments: (a) deterministic response threshold model and (b) extended response threshold model (30 replicates). Teams consisted of 1000 agents, which needed to dynamically self-allocate to two different tasks in order for the team to achieve the optimal performance.

mutations, each allele of the newly added teams to the population was randomly set to a value between  $-1$  to  $+1$  (8-bit encoding, 256 possible values with a resolution of  $\frac{1}{128}$ ) with a probability 0.001. Other settings were the same as in Section III.

**Results:** There were important differences in team performance between the two treatments (RAS and FAS), under both conditions (without and with mutations) and for both models (DTM and ETM). The team performance was higher with FAS than with RAS for both models, when mutations were not used [Fig. 24, top row; mean performance  $\pm$ s.d. at generation 2000 with DTM and without mutations, FAS:  $(7438 \pm 6) \cdot 10^{-4}$ , RAS:  $(6698 \pm 141) \cdot 10^{-4}$ ; with ETM and without mutations, FAS:  $(9161 \pm 9) \cdot 10^{-4}$ , RAS:  $(6695 \pm 225) \cdot 10^{-4}$ ; both Wilcoxon tests,  $df = 29$ ,  $p < 0.001$ ]. With DTM and with mutations, there was a significant difference in performance between the two treatments, however, it was very small [Fig. 24(a), bottom row; mean performance  $\pm$ s.d. at generation 2000, FAS:  $(7304 \pm 7) \cdot 10^{-4}$ , RAS:  $(7298 \pm 10) \cdot 10^{-4}$ ; Wilcoxon test,  $df = 29$ ,  $p < 0.01$ ]. With ETM and with mutations, the performance was higher with FAS than with RAS [Fig. 24(a), bottom row; mean performance  $\pm$ s.d. at generation 2000, FAS:  $(8966 \pm 17) \cdot 10^{-4}$ , RAS:  $(8672 \pm 22) \cdot 10^{-4}$ ; Wilcoxon test,  $df = 29$ ,  $p < 0.001$ ]. The results confirm our observations that using mutations with RAS may to some extent overcome premature convergence (see Appendix C). Nevertheless, for the evolution of more complex agents (here ETM), FAS remained superior to RAS. Moreover, in all cases, FAS led to a faster evolution of teams displaying the highest performance [Fig. 24]. Note that the performance differences between the DTM and ETM are due to limitations in task switching implicitly present in the formulation of the DTM (see [57] for more details).

We chose this setup because it has been used previously [57], [69]. We are aware that every specific problem has its own constraints and limitations. For example, in the scenario

used here, there are two tasks; thus, the optimal team consists probably of a few large groups of agents of the same type, e.g., regulators and foragers. This explains the observed higher performance in treatments with FAS than with RAS (in the main experiment, we showed that using FAS is a more efficient strategy of evolving team compositions when the optimal team has a rather low level of agent heterogeneity). Thus, we believe that the results presented in this Appendix support our earlier conclusions and show that they are also valid for more complex problems.

## APPENDIX H

### FAMILY OF OPERATORS ALTERING TEAM COMPOSITION

Interesting connections between RAS, FAS, and other team composition altering operators exist. First, note that FAS swaps agents on permuted positions, in contrast to RAS which swaps agents on corresponding positions. Consequently, one can implement FAS as a composition of RAS and an operator that permutes the agents' order in a team, i.e., samples agents without repetitions.

A natural next step is to consider a situation when the agents in a team are sampled with repetitions because this allows for a more straightforward implementation of the agents' sampling operator. Let a parenting team have  $M$  agents. In order to construct an offspring team, the agent sampling with repetition (ASWR) picks uniformly at random  $M$  agents with repetitions from the parenting team and puts them in the offspring team. Consider agent  $A$  in the parenting team. The number of its copies in the offspring team follows a binomial distribution with  $M$  trials and the probability of success being  $\frac{1}{M}$ . Thus, if the size of a team is sufficiently large ( $M > 20$ ) one can approximate the binomial distribution by Poisson distribution with the coefficient  $\lambda = 1$ . Consequently, after applying ASWR on a parenting team, each agent will be present in the offspring team  $k$  times with probability  $\frac{1}{e \cdot k!}$ , where  $k = 0, 1, 2, \dots, M$ . Thus, ASWR implements a sort of simultaneous agents' deletion and duplication (or rather multiplication).

Interestingly, operators that delete and duplicate parts of genomes have been already proposed. In genetic programming, there is branch deletion and duplication [33], and in gene expression programming, there is gene transfer [34]. Gene transfer is equivalent to branch substitution, i.e., having two branches  $A$  and  $B$ , delete  $A$  and duplicate  $B$  in place of  $A$ . In the context of evolving team compositions, an analogous agent transfer operator might be proposed. It should have a similar impact on the evolutionary dynamics as FAS. This is because agent transfer would be in essence similar to agent sampling, which, in turn, is a subcomponent of FAS. Note, however, that agent duplication, deletion, and transfer are applied to a single parenting team. Consequently, they should be considered as mutations rather than crossover.

### ACKNOWLEDGMENT

The authors would like to thank A. Maesani, P. Ramdya, and K. Krawiec for their useful comments on this paper.

## REFERENCES

- [1] E. Bonabeau, M. Dorigo, and G. Theraulaz, *Swarm Intelligence: From Natural to Artificial Systems*. New York: Oxford Univ. Press, 1999.
- [2] E. Bonabeau, M. Dorigo, and G. Theraulaz, "Inspiration for optimization from social insect behaviour," *Nature*, vol. 406, no. 6791, pp. 39–42, 2000.
- [3] C. Reynolds, "An evolved, vision-based behavioral model of coordinated group motion," in *Proc. 2nd Int. Conf. Simulat. Adap. Behavior*, 1993, pp. 384–392.
- [4] W. Xiang and H. Lee, "Ant colony intelligence in multi-agent dynamic manufacturing scheduling," *Eng. Applicat. Artif. Intell.*, vol. 21, no. 1, pp. 73–85, 2008.
- [5] G. Mathews, H. Durrant-Whyte, and M. Prokopenko, "Decentralised decision making in heterogeneous teams using anonymous optimisation," *Robot. Autonomous Syst.*, vol. 57, no. 3, pp. 310–320, 2009.
- [6] M. Quinn, L. Smith, G. Mayley, and P. Husbands, "Evolving controllers for a homogeneous system of physical robots: Structured cooperation with minimal sensors," *Philosophic. Trans. Royal Soc. London. Series A: Math., Phys. Eng. Sci.*, vol. 361, no. 1811, pp. 2321–2343, 2003.
- [7] M. Dorigo, V. Trianni, E. Şahin, R. Grob, T. Labella, G. Baldassarre, S. Nolfi, J. Deneubourg, F. Mondada, D. Floreano, L. M. Gambardella, "Evolving self-organizing behaviors for a swarm-bot," *Autonomous Robots*, vol. 17, no. 2, pp. 223–245, 2004.
- [8] C. Parker and H. Zhang, "Cooperative decision-making in decentralized multiple-robot systems: The best-of-n problem," *IEEE/ASME Trans. Mech.*, vol. 14, no. 2, pp. 240–251, Apr. 2009.
- [9] G. Baldassarre, S. Nolfi, and D. Parisi, "Evolving mobile robots able to display collective behaviors," *Artif. Life*, vol. 9, no. 3, pp. 255–267, 2003.
- [10] S. Nolfi and D. Floreano, *Evolutionary Robotics: The Biology, Intelligence, and Technology*. Cambridge, MA: MIT Press, 2000.
- [11] H. Iba, "Emergent cooperation for multiple agents using genetic programming," in *Proc. 4th Int. Conf. Parallel Problem Solv. Nat.*, 1996, pp. 32–41.
- [12] L. Panait and S. Luke, "Cooperative multi-agent learning: The state of the art," *Autonomous Agents Multi-Agent Syst.*, vol. 11, no. 3, pp. 387–434, 2005.
- [13] M. Waibel, L. Keller, and D. Floreano, "Genetic team composition and level of selection in the evolution of cooperation," *IEEE Trans. Evol. Computat.*, vol. 13, no. 3, pp. 648–660, Jun. 2009.
- [14] T. Miconi, "When evolving populations is better than coevolving individuals: The blind mice problem," in *Proc. 18th Int. Joint Conf. Artif. Intell.*, 2003, pp. 647–652.
- [15] J. Bongard, "The legion system: A novel approach to evolving heterogeneity for collective problem solving," in *Proc. EuroGP*, 2000, pp. 16–28.
- [16] S. Luke and L. Spector, "Evolving teamwork and coordination with genetic programming," in *Proc. 1st Annu. Conf. Genet. Program.*, 1996, pp. 150–156.
- [17] T. Haynes and S. Sen, "Crossover operators for evolving a team," in *Proc. 2nd Annu. Conf. Genet. Program.*, 1997, pp. 162–167.
- [18] Y. Suzuki and T. Arita, "A comprehensive evaluation of the methods for evolving a cooperative team," *Artif. Life Robot.*, vol. 10, no. 2, pp. 157–161, 2006.
- [19] K. De Jong, "Learning with genetic algorithms: An overview," *Mach. Learn.*, vol. 3, nos. 2–3, pp. 121–138, 1988.
- [20] M. Brameier and W. Banzhaf, "Evolving teams of predictors with linear genetic programming," *Genet. Program. Evol. Mach.*, vol. 2, no. 4, pp. 381–407, 2001.
- [21] T. Haynes and S. Sen, "Co-adaptation in a team," *Int. J. Computat. Intell. Organiz.*, vol. 1, no. 4, pp. 1–20, 1996.
- [22] D. Andre and A. Teller, "Evolving team Darwin united," in *Proc. RoboCup-98: Robot Soccer World Cup II*, 1999, pp. 346–351.
- [23] S. Luke, C. Hohn, J. Farris, G. Jackson, and J. Hendler, "Co-evolving soccer softbot team coordination with genetic programming," in *Proc. RoboCup-97: Robot Soccer World Cup I*, 1998, pp. 398–411.
- [24] L. Bull, "Evolutionary computing in multi-agent environments: Operators," in *Proc. 7th Int. Conf. EP*, 1998, pp. 370–377.
- [25] A. Eiben, G. Nitschke, and M. Schut, "Collective specialization for evolutionary design of a multi-robot system," in *Proc. Swarm Robot.*, 2007, pp. 189–205.
- [26] D. Floreano, S. Mitri, S. Magnenat, and L. Keller, "Evolutionary conditions for the emergence of communication in robots," *Curr. Biol.*, vol. 17, no. 6, pp. 514–519, 2007.

- [27] T. Soule and R. Heckendorn, "Improving performance and cooperation in multi-agent systems," in *Proc. Genet. Program. Theory Practice V*, 2008, pp. 221–237.
- [28] G. Syswerda, "Uniform crossover in genetic algorithms," in *Proc. 3rd ICGA*, 1989, pp. 2–9.
- [29] L. Eshelman, R. Caruana, and J. Schaffer, "Biases in the crossover landscape," in *Proc. 3rd ICGA*, 1989, pp. 10–19.
- [30] K. De Jong and W. Spears, "A formal analysis of the role of multi-point crossover in genetic algorithms," *Ann. Math. Artif. Intell.*, vol. 5, no. 1, pp. 1–26, 1992.
- [31] T. Nakashima, M. Takatani, M. Udo, H. Ishibuchi, and M. Nii, "Performance evaluation of an evolutionary method for robocup soccer strategies," in *Proc. RoboCup 2005: Robot Soccer World Cup IX*, 2006, pp. 616–623.
- [32] A. Hara and T. Nagao, "Emergence of cooperative behavior using ADG: Automatically Defined Groups," in *Proc. GECCO*, 1999, pp. 1038–1046.
- [33] J. Koza, "Evolving the architecture of a multi-part program in genetic programming using architecture-altering operations," in *Proc. 4th Annu. Conf. Evol. Program.*, 1995, pp. 695–718.
- [34] C. Ferreira, "Gene expression programming: A new adaptive algorithm for solving problems," *Complex Syst.*, vol. 13, no. 2, pp. 87–129, 2001.
- [35] M. Sebag and M. Schoenauer, "Mutation by imitation in boolean evolution strategies," in *Proc. 4th Conf. PPSN*, 1996, pp. 356–365.
- [36] S. Berman, A. Halasz, M. Hsieh, and V. Kumar, "Optimized stochastic policies for task allocation in swarms of robots," *IEEE Trans. Robot.*, vol. 25, no. 4, pp. 927–937, Aug. 2009.
- [37] T. Dahl, M. Mataric, and G. Sukhatme, "Multi-robot task allocation through vacancy chain scheduling," *Robot. Autonomous Syst.*, vol. 57, nos. 6–7, pp. 674–687, 2009.
- [38] S. Beshers and J. Fewell, "Models of division of labor in social insects," *Annu. Rev. Entomol.*, vol. 46, no. 1, pp. 413–440, 2001.
- [39] C. Smith, A. Toth, A. Suarez, and G. Robinson, "Genetic and genomic analyses of the division of labour in insect societies," *Nat. Rev. Genet.*, vol. 9, no. 10, pp. 735–748, 2008.
- [40] A. Duarte, F. Weissing, I. Pen, and L. Keller, "An evolutionary perspective on self-organized division of labor in social insects," *Annu. Rev. Ecol. Evol. System.*, vol. 42, pp. 91–110, Dec. 2011.
- [41] G. Robinson, "Regulation of division of labor in insect societies," *Annu. Rev. Entomol.*, vol. 37, no. 1, pp. 637–665, 1992.
- [42] D. Gordon, "The organization of work in social insect colonies," *Nature*, vol. 380, no. 6570, pp. 121–124, 1996.
- [43] J. Fewell, "Social insect networks," *Science*, vol. 301, no. 5641, p. 1867, 2003.
- [44] E. Robinson, O. Feinerman, and N. Franks, "Flexible task allocation and the organization of work in ants," *Proc. Royal Soc. B: Biol. Sci.*, vol. 276, no. 1677, p. 4373, 2009.
- [45] S. Luke. (2010). *A User Manual for the ECJ Evolutionary Computation Library* [Online]. Available: <http://www.cs.gmu.edu/~eclab/projects/ecj/docs/manual/manual.pdf>
- [46] R. Brooks, "New approaches to robotics," *Science*, vol. 253, no. 5025, p. 1227, 1991.
- [47] R. Brooks, "Intelligence without representation," *Artif. Intell.*, vol. 47, nos. 1–3, pp. 139–159, 1991.
- [48] T. Balch and R. Arkin, "Behavior-based formation control for multirobot teams," *IEEE Trans. Robot. Automat.*, vol. 14, no. 6, pp. 926–939, Dec. 1998.
- [49] M. Campos, E. Bonabeau, G. Theraulaz, and J. Deneubourg, "Dynamic scheduling and division of labor in social insects," *Adap. Behavior*, vol. 8, no. 2, pp. 83–95, 2000.
- [50] C. Jones and M. Mataric, "Adaptive division of labor in large-scale minimalist multi-robot systems," in *Proc. IEEE/RSS Int. Conf. IROS*, Oct. 2003, vol. 2, pp. 1969–1974.
- [51] M. Mataric, G. Sukhatme, and E. Østergaard, "Multi-robot task allocation in uncertain environments," *Autonomous Robots*, vol. 14, no. 2, pp. 255–263, 2003.
- [52] W. Agassounon and A. Martinoli, "Efficiency and robustness of threshold-based distributed allocation algorithms in multi-agent systems," in *Proc. 1st Int. Joint Conf. Autonomous Agents Multiagent Syst.: Part 3*, 2002, pp. 1090–1097.
- [53] P. Lichocki, K. Krawiec, and W. Jaśkowski, "Evolving teams of cooperating agents for real-time strategy game," in *Proc. EvoWorkshops Appl. Evol. Comput.*, 2009, pp. 333–342.
- [54] M. Preuss, N. Beume, H. Danielsiek, T. Hein, B. Naujoks, N. Piatkowski, R. Ster, A. Thom, and S. Wessing, "Towards intelligent team composition and maneuvering in real-time strategy games," *IEEE Trans. Computat. Intell. AI Games*, vol. 2, no. 2, pp. 82–98, Jun. 2010.
- [55] M. Dorigo, M. Birattari, and T. Stutzle, "Ant colony optimization," *IEEE Computat. Intell. Mag.*, vol. 1, no. 4, pp. 28–39, Nov. 2006.
- [56] T. White, B. Pagurek, and F. Oppacher, "ASGA: Improving the ant system by integration with genetic algorithms," in *Proc. 3rd Annu. Conf. Genet. Program.*, 1998, pp. 610–617.
- [57] P. Lichocki, D. Tarapore, L. Keller, and D. Floreano, "Neural networks as mechanisms to regulate division of labor," *Am. Naturalist*, vol. 179, no. 3, pp. 391–400, 2012.
- [58] S. Goings, J. Clune, C. Ofria, and R. Pennock, "Kin selection: The rise and fall of kin-cheaters," in *Proc. 9th Int. Conf. Simul. Syn. Living Syst. Artif. Life*, 2004, pp. 303–308.
- [59] S. Mitri, D. Floreano, and L. Keller, "The evolution of information suppression in communicating robots with conflicting interests," *Proc. Nat. Acad. Sci.*, vol. 106, no. 37, pp. 15786–15790, 2009.
- [60] M. Waibel, D. Floreano, and L. Keller, "A quantitative test of Hamilton's rule for the evolution of altruism," *PLoS Biol.*, vol. 9, no. 5, p. e1000615, 2011.
- [61] Y. Shoham and K. Leyton-Brown, *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge, MA: Cambridge Univ. Press, 2009.
- [62] M. Waibel, D. Floreano, S. Magnenat, and L. Keller, "Division of labour and colony efficiency in social insects: Effects of interactions between genetic architecture, colony kin structure and rate of perturbations," *Proc. Royal Soc. B: Biol. Sci.*, vol. 273, no. 1595, pp. 1815–1823, 2006.
- [63] D. Goldberg, K. Deb, and J. Clark, "Genetic algorithms, noise, and the sizing of populations," *Complex Syst.*, vol. 6, no. 4, pp. 333–362, 1991.
- [64] K. De Jong, "Genetic algorithms are not function optimizers," in *Proc. 2nd Workshop Found. Genet. Algorithms*, 1993, pp. 5–17.
- [65] L. Whitley, "Fundamental principles of deception in genetic search," in *Proc. 1st Workshop Found. Genet. Algorithms*, 1991, pp. 221–241.
- [66] J. Horn and D. Goldberg, "Genetic algorithm difficulty and the modality of fitness landscapes," in *Proc. 3rd Workshop Found. Genet. Algorithms*, 1994, pp. 243–269.
- [67] R. Page, Jr., and S. Mitchell, "Self-organization and the evolution of division of labor," *Apidologie*, vol. 29, nos. 1–2, pp. 171–190, 1998.
- [68] E. Bonabeau, G. Theraulaz, and J. Deneubourg, "Quantitative study of the fixed threshold model for the regulation of division of labour in insect societies," *Proc. Biol. Sci.*, vol. 263, no. 1376, pp. 1565–1569, 1996.
- [69] D. Tarapore, D. Floreano, and L. Keller, "Task-dependent influence of genetic architecture and mating frequency on division of labour in social insect societies," *Behavioral Ecol. Sociobiol.*, vol. 64, no. 4, pp. 675–684, 2010.
- [70] T. Miconi, "A collective genetic algorithm," in *Proc. GECCO*, 2001, pp. 876–883.
- [71] G. Werner and M. Dyer, "Evolution of communication in artificial organisms," in *Proc. 2nd Int. Conf. Artif. Life*, 1991, pp. 659–687.
- [72] S. Ficici, R. Watson, and J. Pollack, "Embodied evolution: A response to challenges in evolutionary robotics," in *Proc. 8th Eur. Workshop Learn. Robots*, 1999, pp. 14–22.
- [73] K. Stanley, B. Bryant, and R. Miikkulainen, "Real-time neuroevolution in the nero video game," *IEEE Trans. Evol. Computat.*, vol. 9, no. 6, pp. 653–668, Dec. 2005.
- [74] R. Watson, S. Ficici, and J. Pollack, "Embodied evolution: Distributing an evolutionary algorithm in a population of robots," *Robot. Autonomous Syst.*, vol. 39, no. 1, pp. 1–18, 2002.
- [75] L. Spector, J. Klein, C. Perry, and M. Feinstein, "Emergence of collective behavior in evolving populations of flying agents," *Genet. Program. Evol. Mach.*, vol. 6, no. 1, pp. 111–125, 2005.
- [76] C. Ward, F. Gobet, and G. Kendall, "Evolving collective behavior in an artificial ecology," *Artif. Life*, vol. 7, no. 2, pp. 191–209, 2001.
- [77] E. Simoes and D. Barone, "Predation: An approach to improving the evolution of real robots with a distributed evolutionary controller," in *Proc. IEEE ICRA*, Aug. 2002, vol. 1, pp. 664–669.
- [78] R. Bianco and S. Nolfi, "Toward open-ended evolutionary robotics: Evolving elementary robotic units able to self-assemble and self-reproduce," *Connect. Sci.*, vol. 16, no. 4, pp. 227–248, 2004.
- [79] A. Agah and G. Bekey, "Phylogenetic and ontogenetic learning in a colony of interacting robots," *Autonomous Robots*, vol. 4, no. 1, pp. 85–100, 1997.
- [80] A. Cangelosi and D. Parisi, "The emergence of a 'language' in an evolving population of neural networks," *Connect. Sci.*, vol. 10, no. 2, pp. 83–97, 1998.



- [81] M. Quinn, "A comparison of approaches to the evolution of homogeneous multi-robot teams," in *Proc. CEC*, 2001, vol. 1, pp. 128–135.
- [82] A. Agah and K. Tanie, "Robots playing to win: Evolutionary soccer strategies," in *Proc. IEEE ICRA*, Apr. 1997, vol. 1, pp. 632–637.
- [83] M. Potter and K. De Jong, "A cooperative coevolutionary approach to function optimization," in *Proc. 3rd Int. Conf. PPSN*, 1994, pp. 249–257.
- [84] M. A. Potter and D. V. De Jong, "Cooperative coevolution: An architecture for evolving coadapted subcomponents," *Evol. Computat.*, vol. 8, no. 1, pp. 1–29, 2000.
- [85] R. Wiegand, W. Liles, and K. De Jong, "An empirical analysis of collaboration methods in cooperative coevolutionary algorithms," in *Proc. GECCO*, 2001, pp. 1235–1242.
- [86] M. Mirolli and D. Parisi, "How can we explain the emergence of a language that benefits the hearer but not the speaker?" *Connect. Sci.*, vol. 17, no. 3, pp. 307–324, 2005.
- [87] L. Panait, S. Luke, and R. Wiegand, "Biasing coevolutionary search for optimal multiagent behaviors," *IEEE Trans. Evol. Computat.*, vol. 10, no. 6, pp. 629–645, Dec. 2006.
- [88] H. Botee and E. Bonabeau, "Evolving ant colony optimization," *Complex Syst.*, vol. 1, no. 2, pp. 149–159, 1998.
- [89] A. Agogino, K. Stanley, and R. Miikkulainen, "Online interactive neuro-evolution," *Neural Process. Lett.*, vol. 11, no. 1, pp. 29–38, 2000.
- [90] J. Grefenstette, "Credit assignment in rule discovery systems based on genetic algorithms," *Mach. Learn.*, vol. 3, no. 2, pp. 225–245, 1988.



**Paweł Lichocki** (M'07) received the B.S. and M.S. degrees in computer science from the Poznan University of Technology, Poznan, Poland, with specialization in intelligent decisions support systems, in 2007 and 2008, respectively. He is currently pursuing the Ph.D. degree in computer, communication, and information sciences.

He is currently a Doctoral Assistant with the Laboratory of Intelligent Systems, School of Engineering, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland. His current research interests

include evolution of natural and artificial societies (social insects, teams of robots, etc.) and bioinspired control of multiagent systems.



**Steffen Wischmann** received the M.S. degree in biology from Saarland University, Saarbrücken, Germany, in 2003, and the Ph.D. degree in theoretical biology from the University of Bonn, Bonn, Germany, in 2008.

He is currently a joint Post-Doctoral Fellow with the Laboratory of Intelligent Systems, School of Engineering, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland, and with the Department of Ecology and Evolution, University of Lausanne, Lausanne. His current research interests

include proximate and ultimate mechanisms driving the evolution of intra and interspecific social behavior, evolutionary robotics, and information theoretic approaches to understand communication and cooperation.



**Laurent Keller** received the B.S. and M.S. degrees in biology and the Ph.D. degree in zoology from the University of Lausanne, Lausanne, Switzerland, in 1983, 1985, and 1989, respectively.

He is currently a Professor of Ecology and Evolution and the Head of the Department of Ecology and Evolution, Biophore, University of Lausanne. His current research interests include the principles governing the evolution of animal societies and the ecological and evolutionary consequences of social life. In addition to publishing several research papers, he has edited two books: the *Queen Number and Sociality in Insects*, in 1993, and the *Levels of Selection in Evolution*, in 1999.

Dr. Keller received the E. O. Wilson Naturalist Award in 2005.



**Dario Floreano** (SM'05) received the M.A. degree in visual psychophysics from the University of Trieste, Trieste, Italy, in 1988, the M.S. degree in neural computation from the University of Stirling, Stirling, U.K., in 1992, and the Ph.D. degree in cognitive systems and robotics from the University of Trieste in 1995.

He is currently an Associate Professor and the Director of the Laboratory of Intelligent Systems, School of Engineering, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland. His current research interests include analysis and synthesis of bioinspired adaptive systems, such as biomimetic robotics, neuromorphic engineering, and artificial evolution. He is the co-author of two books on these topics: the *Evolutionary Robotics* (Cambridge, MA: MIT Press, 2000) and the *Bio-Inspired Artificial Intelligence* (Cambridge, MA: MIT Press, 2008).