

The Evolution of Division of Labor

Heather J. Goldsby, David B. Knoester, Jeff Clune,
Philip K. McKinley, Charles Ofria

Department of Computer Science and Engineering
Michigan State University
East Lansing, Michigan 48824
{hjb, dk, jclune, mckinley, ofria}@cse.msu.edu

Abstract. We use digital evolution to study the division of labor among heterogeneous organisms under multiple levels of selection. Although division of labor is practiced by many social organisms, the labor roles are typically associated with different individual fitness effects. This fitness variation raises the question of why an individual organism would select a less desirable role. For this study, we provide organisms with varying rewards for labor roles and impose a group-level pressure for division of labor. We demonstrate that a group selection pressure acting on a heterogeneous population is sufficient to ensure role diversity regardless of individual selection pressures, be they beneficial or detrimental.

Keywords: digital evolution, cooperative behavior, specialization, altruism.

1 Introduction

Within nature, many organisms live in groups where individuals assume different roles and cooperate to survive [1–4]. For example, in honeybee colonies, among other roles, drones care for the brood, workers forage for pollen, and the queen focuses on reproduction [1]. A notable aspect of these roles is that they do not all accrue the same fitness benefits. For example, leadership is a common role found within multiple species, where the benefits of leadership are significantly greater than that of a follower. In human societies, leaders of organizations commonly earn many times more than the average worker [4]. An open question is why individuals would not all attempt to perform the role associated with the highest fitness benefit, or in other words, why individuals would perform roles that put their genes at an evolutionary disadvantage for survival. Group selection pressures among human tribes have been proposed as one possible explanation for the evolution of leaders and followers [4]. In this paper, we explore whether group selection is sufficient to produce division of labor, where individual selection rewards different roles unequally.

Numerous evolutionary computation approaches have been used to study the behavior of cooperative groups comprising heterogeneous members [5–9]. Two key differentiating characteristics for these approaches are the *level of selection* used (i.e., individual or group) and whether or not *division of labor*

occurs. Ecological approaches [6] use individual-level selection in concert with limited resources to promote the evolution of specialists. Some coevolutionary approaches [5, 9] evolve cooperative groups using individual selection, where different species are isolated in distinct subpopulations. Cooperation among these species occurs only at the time of fitness evaluation, when individuals from one species are evaluated with representatives from each of the other species. Perez-Uribe *et al.* [7] and Waibel *et al.* [8] overview work performed in this area, and also describe the effects of varying the levels of selection and population composition (i.e., heterogeneous or homogeneous populations) [8]. However, these prior studies do not address multi-level selection, where organisms experience individual-level competition to survive and also group-level pressure to cooperate. To explore these conditions, which are pertinent to biological studies of the division of labor in cooperative groups, we apply multi-level selection to a heterogeneous population.

For this study, we use AVIDA [10], a digital-evolution platform previously used to study topics including the origin of complex features [11] and the evolution of cooperation among homogeneous individuals [12]. Within an AVIDA experiment, a population of self-replicating computer programs exists in a user-defined computational environment and is subject to mutations and natural selection. These digital organisms execute their genome to perform tasks that metabolize resources in the environment, interact with neighboring organisms, and self-replicate.

In this paper, we describe how we used AVIDA and multi-level selection to evolve groups of heterogeneous organisms that perform roles with different fitness benefits. First, we enabled organisms to self-select roles associated with different costs and/or benefits. We applied a group-level pressure for division of labor that successfully counteracted the individual pressure to perform only the highest rewarded role. Second, rather than having an organism select a role, we conducted experiments in which a role was associated with a labor task that the organism had to perform. Again, we observed the evolution of division of labor. Third, we analyzed one of the successful groups and determined that it used a combination of genotypic diversity, phenotypic plasticity, and cooperation to perform all roles. The model developed for this approach can be used to inform biological studies of cooperation, such as those performed by Dornhaus *et al.* for honeybees [1]. Additionally, this technique can serve as a means to achieve division of labor within artificial life in order to solve engineering problems, such as developing multiple software components that must interact to achieve an overall objective [12], as well as cooperation among heterogeneous robots [13].

2 Methods

For each experiment, 30 trials were conducted to account for the stochastic nature of evolution. Figure 1 depicts an AVIDA organism and population. An AVIDA organism consists of a circular list of instructions (its *genome*) and a virtual CPU that executes those instructions. The virtual CPU architecture comprises three general-purpose registers $\{AX, BX, CX\}$ and two stacks

$\{GS, LS\}$. The standard AVIDA instruction set used in this study is Turing complete and is designed so that random mutations will always yield a syntactically correct program, albeit one that may not perform any meaningful computation [14]. This AVIDA instruction set performs basic computational tasks (addition, multiplication, and bit-shifts), controls execution flow, enables communication, and allows for replication. In this study, the instruction set also included several instructions developed for the evolution of distributed problem solving [12]; these instructions are summarized in Table 1.

AVIDA organisms can perform *tasks* that enable them to metabolize resources from their environment. It is typical for these tasks to be logical operations performed on 32-bit integers. Performing tasks increases an organism’s *merit*, which determines the rate at which its virtual CPU will execute instructions relative to the other organisms in the population. For example, an organism with a merit of 2 will, on average, execute twice as many instructions as an organism with a merit of 1. Since organisms self-replicate, an organism with greater merit will generally out-compete other organisms, eventually dominating the population. For these experiments, the amount of merit that an organism gains for completing a task depends on a user-defined constant called the task’s *bonus value*. When an organism performs a task, the organism’s merit is multiplied by the task’s bonus value. For example, if an organism performs a task with a bonus value of 2, its merit is doubled.

An AVIDA population comprises a number of *cells* in which at most one organism can live. Thus, the size of an AVIDA population is bounded by the number of cells in the environment. The cells are divided into a set of distinct subpopulations, called *demes*. In this study, demes compete every 100 *updates* in a tournament based on their fitness function, where a deme’s fitness is determined by the behavior of its constituent organisms. An update is the unit of experimental time in AVIDA corresponding to approximately 30 instructions per organism. Each tournament contains a set of demes selected at random, and the deme with greatest fitness replaces the other demes (ties are broken randomly). When a deme is replaced, all organisms from the source deme are copied into the target deme, overwriting its previous inhabitants. Within each deme, organisms are still able to self-replicate. Thus, an individual’s survival is dependent not only on its ability to out-compete its neighbors for the limited space available in its deme, but also on the collective behavior of the group. This process is similar to competition among human tribes [3].

For the experiments described in this paper, we created mutually-exclusive tasks for each possible role. An organism fulfills a role by performing its associated task; each organism may only have one role. For some of our experiments, we




Fig. 1. An Avida organism and population.

Table 1. Communication and coordination instructions for this study.

Instruction	Description
send-msg	Sends a message to the neighbor currently faced by the caller.
retrieve-msg	Loads the caller’s virtual CPU from a previously received message.
rotate-left-one	Rotate this organism counter-clockwise one step.
rotate-right-one	Rotate this organism clockwise one step.
get-role-id	Sets register BX to the value of the caller’s <code>role-id</code> register.
set-role-id	Sets the caller’s <code>role-id</code> register to the value in register BX .
bcast1	Sends a message to all neighboring organisms.
get-cell-xy	Sets register BX and CX to the $x - y$ coordinates of the caller.
collect-cell-data	Sets register BX to the value of the cell data where the caller lives.

used *role-ids*, a mechanism whereby an organism sets a special-purpose virtual CPU register to an integer value, to indicate the role that an organism performs. For others, we required the organisms to implement logical operations. We tried these two mechanisms to see if the complexity of labor tasks affected the division of labor. We varied the benefits of performing a task (and thus performing a role) by changing the task’s bonus value. In all experiments presented here, we used 400 demes, each containing 25 organisms on a 5×5 toroidal grid. Deme fitness was based on the diversity of tasks performed by the organisms. Thus, our experiments contain both the individual-level pressure to perform the task with the highest reward and the group-level pressure to diversify the tasks performed.

3 Experimental Results

Varying Roles. Our first experiment was designed to ascertain whether group selection is a strong enough pressure to produce division of labor when all roles have the same fitness benefit. For this experiment, we considered an organism to be performing a given role if it sets its `role-id` register to a desired value using the `set-role-id` instruction. We varied the desired number of roles from 2 to 20 and associated each `role-id` with a task that has a bonus value of 2. For example, when two roles are desired, the rewarded `role-ids` are 1 and 2. If an organism replicates after setting its `role-id` to 1, then it has completed task 1 and as a result, its merit is doubled. Similarly, if five roles are desired, then the rewarded `role-ids` are $\{1, 2, 3, 4, 5\}$. If an organism sets its `role-id` to a value outside this range, then no reward is granted. Additionally, we impose a group-level pressure for both the number of organisms that have set a `role-id` and the diversity of the `role-ids`. Specifically, the deme fitness function used here is:

$$F = \begin{cases} 1 + n & \text{if } n < 25 \\ 1 + n + r & \text{if } n = 25 \end{cases} \quad (1)$$

where F is the fitness of a deme, n is the number of organisms that have set a `role-id`, and r is the number of unique rewarded `role-ids` set by organisms in the deme. Experiments described in this paper were repeated with tournaments of size 2, 5, and 10; results were not qualitatively different. Due to space limitations, we present results for a tournament of size 2, except where noted.

Figure 2 depicts the grand mean and maximum performance of all demes across 30 trials for each of $\{2, 3, 4, 5, 10, 20\}$ roles. The different curves represent the varying number of desired roles. In general, the best performing deme

achieves the desired number of roles within 5,000 to 15,000 updates. Our analysis of the behavior of the demes indicates that they exploit location awareness (using instruction `get-cell-xy`) to determine which role to perform. Specifically, their x (or y) coordinate was used to calculate their role. Thus, we conclude that group selection is strong enough to produce division of labor among equally rewarded roles.

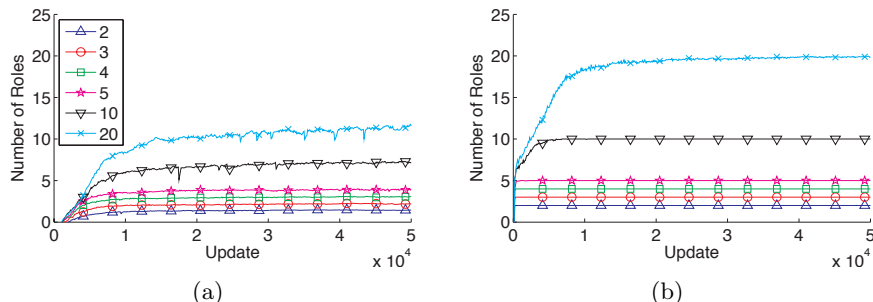


Fig. 2. (a) grand mean and (b) grand maximum number of unique roles over all demes when the number of desired roles was varied from 2 to 20.

Varying Rewards. In the next set of experiments, we explored whether the group selection pressure for division of labor was strong enough to counteract rewarding roles unequally. The different rewards associated with roles provides an individual pressure to specialize on the most rewarded role, even if this behavior is detrimental to the performance of the group. This setup is designed to reflect a leader/follower situation, where it is desirable for the group to have one leader, and yet the rewards for the leader may be significantly different than those of a follower. To test this, we set the desired number of roles to be 2, and conducted trials for different multiplicative benefits of role-id 1 (the leader role). All other role-ids were neutral (i.e., they did not affect merit). We then specified a group-level pressure to limit the leader role to only a single organism. The deme fitness function used here was:

$$F = \begin{cases} 1 + n & \text{if } n < 25 \\ (1 + n - (o_1 - d_{o_1}))^2 & \text{if } n = 25 \end{cases} \quad (2)$$

where F is the fitness of a deme, n is the number of organisms that have set a role-id, d_{o_1} is the desired number of organisms that perform the leader role, and o_1 is the actual number of organisms that perform the leader role.

Figure 3 depicts the results of varying the multiplicative benefit of the leader role from 0.5 (a penalty) to 64 (a significant reward) across 30 trials. Each treatment has two lines: a dashed line to indicate the number of followers, and a solid line to indicate the number of leaders; different symbols are used to indicate different treatments. In general, by 50,000 updates, the average deme has reached an equilibrium between leaders and followers, where the number of leaders is less than five in all treatments. The larger the benefit of leadership, the slower the population was to reach equilibrium and the larger the number of leaders. These

results indicate that group selection is able to effectively counteract individual selection pressures.

To assess the generality of these results, we ran a control experiment without group selection, where the fitness of all demes were always equivalent. As expected, the majority of organisms chose to be leaders when there was a reward and followers when there was a penalty. Additionally, we ran two-role experiments where we set the desired number of leaders to be 13 (approximately half of the population). Similar to the results depicted in Figure 3, the population reached equilibrium by 50,000 updates with the population nearly evenly divided between leaders and followers. Lastly, we conducted treatments where we set the desired number of role-ids to be five and varied the distribution of the benefits among roles. Specifically, we conducted experiments where the benefits: (1) increased linearly; (2) where one role was rewarded significantly greater than the others; and (3), where the majority of task bonus values were penalties. For these experiments, the deme fitness function was the number of unique tasks performed. In all cases, the best performing deme rapidly evolved to perform all five roles, with the average deme performing three or more roles.




Fig. 3. Group selection is strong enough to overcome individual pressures for leadership, which we test by varying the multiplicative bonus value of leadership from 0.5 to 64.

Increasing Complexity. For the last set of experiments, we studied whether this multi-level selection technique was sufficient to evolve division of labor when the complexity of the corresponding tasks varied. In the first two sets of experiments, an organism performed a role by setting its role-id to a specific value. For these experiments, we required the organism to perform a bit-wise logic operation on 32-bit integers. In this case, we used five mutually-exclusive logic operations: `not`, `nand`, `and`, `orn`, or `or`. For the first experiment in this set, we rewarded all five logic operations equally with a task bonus value of 2. The deme fitness function was set to the number of unique tasks performed plus 1. The best demes varied between 4 and 5 tasks, whereas the average deme consistently performed between 3 and 4 tasks. Thus, this technique was successful in producing division of labor.

Next we examined whether division of labor occurred when the tasks were complex and their rewards were unequal. We rewarded the tasks based on complexity. Task `not` was assigned a bonus value of 2, tasks `nand`, `and`, `orn` were assigned a bonus value of 3, and `or` was assigned a bonus value of 4. This treatment increases the difficulty of the problem because organisms must evolve to perform logic operations and coordinate roles with different benefits. The best performing demes continued to vary between 4 and 5 tasks, and the average deme continued to perform between 3 and 4 tasks. This result indicates that

the group selection pressure is strong enough to counteract the varying rewards among complex individual tasks.

Behavioral Analysis. To better understand how group selection maintained diversity of all five logic tasks in the previous treatment, we analyzed one deme from the end of a trial where all tasks had equivalent rewards. This deme had 18 unique genotypes. After we put the deme through a period of 100 updates without mutations, an *ecological period*, the deme maintained all 5 tasks and had 6 different genotypes. Figure 4 provides a graphical depiction of the genotype (shading) and phenotype (task-id of the task performed) of the organisms in the deme. Of the 6 genotypes, five exhibited *phenotypic plasticity*. Specifically, depending on their environmental context, genotypes in these families could perform one of two different tasks, *nand/orn* and *not/and*, respectively. For example, the first two blocks are both shaded white, indicating they have the same genotype, and yet one performs task 4 (*orn*) and the other performs task 2 (*and*). The remaining genotype (depicted in black) exclusively performed task 5 (*or*). It is often the case that explanations from a group selection perspective can also be interpreted from a kin selection perspective, and both are equally valid yet provide different intuitions [15]. For this paper, interpreting the results from a group selection perspective provides the most intuitive explanations. Additionally, our analyses revealed that the different genotypes belonged to three distantly-related lineages, with those performing tasks 1-4 more closely related to each other than to the genotype performing task 5.

Lastly, we conducted “knockout” experiments for each communication (`send-msg`, `retrieve-msg`) and environment-sensing instruction (`get-cell-xy`, `collect-cell-data`). Specifically, all instances of the target instruction were replaced with a placeholder instruction that performs no useful function. We conducted 30 trials for each instruction knockout, and the results indicated that the ability to communicate with neighboring organisms and to sense their environment were critical for the success of the group. Without these instructions, at the end of a 100 update period (that did include mutations), the deme performed an average of 2.2 tasks. In summary, the most effective deme strategy relied on a combination of genotypic diversity, phenotypic plasticity, and cooperation to achieve all five tasks.

4 Conclusion

In this paper, we demonstrated that group selection is a sufficient pressure to produce division of labor among heterogeneous organisms. Specifically, we found that group selection can produce demes whose constituent organisms perform

4	2	4	2	4
3	4	2	4	3
1	2	3	1	4
1	3	5	3	1
2	2	3	5	4

Fig. 4. A deme after an ecological period. Genotypes are denoted by different color shading. Phenotypes are indicated by task-id of the role performed.

five different, mutually-exclusive complex roles, regardless of the underlying reward structure. In future work, we seek to use this technique to better understand behavior of social organisms and to harness this approach to apply evolutionary computation to complicated problems, such as automatically generating software systems [12].

Acknowledgements. The authors thank Ben Beckmann, Jeff Barrick, Rob Pennock, and Carlos Navarrete for their contributions. This work was supported by the U.S. National Science Foundation grants CCF-0820220, CCF-0643952, and CCF-0523449.

References

1. Dornhaus, A., Klügl, F., Puppe, F., Tautz, J.: Task selection in honeybees - experiments using multi-agent simulation. In: Proc. of the Third German Workshop on Artificial Life. (1998)
2. Robinson, G.E.: Regulation of division of labor in insect societies. *Annual Review of Entomology* **37** (1992) 637–665
3. Bowles, S., Hopfensitz, A.: The co-evolution of individual behaviors and social institutions. *Journal of Theoretical Biology* **223**(2) (2003) 135–147
4. VanVugt, M., Hogan, R., Kaiser, R.: Leadership, followership, and evolution: Some lessons from the past. *American Psychologist* **63** (2008) 182–196
5. Potter, M.A., DeJong, K.A.: Cooperative coevolution: An architecture for evolving coadapted subcomponents. *Evol. Comput.* **8**(1) (2000) 1–29
6. Goings, S., Ofria, C.: Ecological approaches to diversity maintenance in evolutionary algorithms. In: IEEE Symposium on Artificial Life, Nashville, TN (2009)
7. Perez-Uribe, A., Floreano, D., Keller, L.: Effects of group composition and levels of selection in the evolution of cooperation in artificial ants. In: European Conference on Artificial Life (ECAL). (2003) 128–137
8. Waibel, M., Keller, L., Floreano, D.: Genetic team composition and level of selection in the evolution of cooperation. *IEEE Transactions on Evolutionary Computation* (2009 (To appear))
9. Yong, C.H., Miikkulainen, R.: Cooperative coevolution of multi-agent systems. Technical report, University of Texas at Austin (2001)
10. Ofria, C., Wilke, C.O.: Avida: A software platform for research in computational evolutionary biology. *Journal of Artificial Life* **10** (2004) 191–229
11. Lenski, R.E., Ofria, C., Pennock, R.T., Adami, C.: The evolutionary origin of complex features. *Nature* **423** (2003) 139–144
12. McKinley, P.K., Cheng, B.H.C., Ofria, C., Knoester, D.B., Beckmann, B., Goldsby, H.J.: Harnessing digital evolution. *IEEE Computer* **41**(1) (2008)
13. Labella, T.H., Dorigo, M., Deneubourg, J.L.: Division of labor in a group of robots inspired by ants foraging behavior. *ACM Transactions on Autonomous and Adaptive Systems (TAAS)* (2006)
14. Ofria, C., Adami, C., Collier, T.C.: Design of evolvable computer languages. *IEEE Transactions in Evolutionary Computation* **6** (2002) 420–424
15. Foster, K., Wenseleers, T., Ratnieks, F., Queller, D.: There is nothing wrong with inclusive fitness. *Trends in Ecology and Evolution* **21** (2006) 599–600