Evolutionary Transition through a New Multilevel Selection Model

Shelly Xiaonan Wu[1] and Wolfgang Banzhaf[1]

[1]Department of Computer Science, Memorial University, St John's, NL, Canada A1B 5X3
{xiaonan, banzhaf}@mun.ca

Abstract

Most multilevel selection models in the literature focus on addressing the evolution of cooperation. There is, however, another aspect of multilevel selection theory. It might be able to provide explanations for evolutionary transitions, which involve the creation of higher level complexes out of simpler elements. Here, we propose a multilevel selection model to support evolutionary transitions. This model employs a genetic operator called "cooperation" to build the hierarchical structure used in multilevel selection theory, and applies two types of multilevel selection to achieve transitions. Our experiments on an extended N-player Prisoner's Dilemma game demonstrate that groups with all required skills emerge from a population of independent individuals, no matter whether skills are equally rewarded or not. Our experiments confirm that both types of multilevel selection mentioned are relevant to evolutionary transitions.

Introduction

Our biological world is hierarchically organized. Starting from the bottom level to the top, the hierarchy includes atoms, molecules, organelles, cells, tissues, organs, organ systems, organisms, populations, communities, ecosystems and biospheres. It is also generally accepted that the simpler, smaller components appeared before the more complex, composite systems. The creation of new higher level complexes out of simpler entities is referred to as an "evolutionary transition" (Buss, 1987; Michod, 1999; Smith and Szathmáry, 1995).

How and why evolutionary transitions take place during evolution is an important question to address for biologists and sociologists. Increasingly, multilevel selection (MLS) has been suggested as a potent explanation (Michod, 1999; Smith and Szathmáry, 1995; Sober and Wilson, 1999). MLS theory posits that natural selection may simultaneously operate at multiple levels of the biological hierarchy. Multilevel selection theory has its origins in group selection theory, which initially was aimed to explain the evolution of cooperation[1]: Individuals are divided into groups; within-group selection favors selfish individuals, while between-group selection favors cooperative individuals. When between-group selection dominates within-group selection, a major transition occurs and the group becomes a higher level organism in its own right (Wilson and Wilson, 2007).

The way to explain evolutionary transitions extends MLS theory in an important new way. Nevertheless, investigations of most existing MLS models focus on the conditions necessary for the emergence of cooperation during evolution. The purpose of this paper is to computationally verify the idea that evolutionary transitions can indeed occur through multilevel selection. To this end, we consider a new MLS model and investigate its ability to exploit the division of labor. A crucial step in many of the major transitions (Smith and Szathmáry, 1995) is the division of labor between components of an emerging higher level unit of evolution (Gavrilets, 2010). This new MLS model distinguishes itself from existing MLS models in two ways. First, it integrates two types of multilevel selection (Okasha, 2005), which are believed to be relevant to the evolutionary transitions, each at a different stage. To encourage a transition, group fitness (fitness of higher level units) is defined to be "decoupled" (Michod and Nedelcu, 2003) from the individual fitness (fitness of the lower level units). Second, the model does not take the existence of the hierarchical structure for granted; multicellular organisms do not exist at the beginning of life. Our model constructs the hierarchy through evolutionary transitions. The experiments shown here confirm that in appropriately defined models independent individuals are able to transit to groups with totally different functionalities using multilevel selection; in terms of the division of labor, those are groups with members executing various skills with possibly different rewards.

The remainder of this paper is organized as follows. Section 2 briefly describes multilevel selection theory, especially the two types of multilevel selection. Section 3 in-

---

[1] Group selection is a longstanding controversial area in the evolution of cooperation. It recently re-emerged as an important component of a multilevel theory of evolution. Many strong advocates of other alternatives in explaining the evolution of cooperation have come to accept multilevel analysis (Borrello, 2005; Okasha, 2001, 2008; Wilson, 1983).

troduces our multilevel selection model. Section 4 shows experiments with the model and their results. Section 5 concludes and discusses future work.

Multilevel Selection

Group selection (Sober and Wilson, 1999) tries to explain the evolution of cooperation by introducing selection between groups. Between-group competition allows traits to arise from evolution that are costly for individuals but beneficial to groups. This is therefore one mechanism by which cooperation is able to emerge in evolution. Individuals and groups, however, are relative: an entity can be regarded as a group for individuals at the level below, and as an individual of a group at the level above. This new perspective is now called multilevel selection (MLS) theory.

When higher level selection (i.e. between-group selection) dominates lower level selection (i.e. within-group selection), an evolutionary transition occurs (Wilson and Wilson, 2007). The reason that individuals would give up their survival and reproductive opportunity to become a part of complexes is that the complexes are able to protect their members from being eliminated by selection. For example, by hunting together or by watching predators for others, members in a group have a greater chance to survive severe competition. In addition, a consequence of higher level selection is adaptation, which minimizes conflict among lower level entities and increases cooperation. Therefore, lower level selection does not interrupt the formation of higher level entities (Okasha, 2005).

For the hierarchical structure used in MLS with a number of individual entities nested within each group entity, we need to clarify which entities should become the objects of evolution or which level should undergo evolution (Okasha, 2005). If we are interested in the changing frequencies of different individual traits, individual entities will be the objects of evolution; group entities are only a structure or an environment where fitness-affecting interactions take place. Most multilevel selection models proposed for the evolution of cooperation, such as Wilson (1975)'s and Traulsen and Nowak (2006)'s models, belong to this kind. These models focus on how to propagate the altruistic trait among individuals in a population. To this end, groups are regularly formed and evaluated. Groups with more altruists will have a higher fitness; hence cooperative individuals in such groups will have higher probabilities to be reproduced. In other words, groups are only temporary fitness-bearing entities; even though they are selected, it is not them but individuals that are reproduced, and also it is the frequency of individual traits that is changed. This type of MLS is called MLS type 1 (MLS1) (Damuth and Heisler, 1988; Okasha, 2005).

Alternatively, if we are interested in the changing frequencies of different group traits, group entities need to be the objects of evolution. They are not merely an environment to individual entities or an object of selection; they actually have their own heritable traits. Group entities with higher fitness will reproduce more offspring group entities with similar traits. Individual entities may still undergo evolution within each group entity, which leads to changes in the distribution of individual traits and potentially affects group traits. This type of MLS is called MLS type 2 (MLS2) (Damuth and Heisler, 1988; Okasha, 2005). As a result, since the entities undergoing evolution are different in these two types of multilevel selection, the evolutionary changes obtained on each level are different. MLS1 will contribute the most individual entities to the next generation, while MLS2 will contribute the most groups. Both MLS1 and MLS2 are distinct processes that can occur in nature.

According to Okasha (2005), both types of multilevel selection may be relevant to evolutionary transitions. An evolutionary transition is more complicated than the evolution of cooperation. However, before transitions take place and complexes emerge, simpler entities which constitute the complexes have to be able to work together. They need to sacrifice their individuality and exhibit cooperative traits. Therefore, in the early stage of evolutionary transitions, the evolution of cooperation has to emerge, so that cooperative traits can spread among simpler entities in the population. That is exactly what MLS1 promotes: using groups as an environment to help individual traits to propagate. Once individuals are willing to form cohesive complexes, evolution should work on complexes to gradually develop their own traits. In other words, complexes should now themselves become objects of evolution. Through selection and reproduction, complexes are better adapted to their environment and eventually become discrete units, normally with traits different from their constituents' traits. It follows that MLS2 should be applied at a later stage of an evolutionary transition.

The shift from MLS1 to MLS2 also indicates a change in the definition of group fitness. In MLS1, group fitness is defined as the average fitness of the individuals within a group, while in MLS2, group fitness is defined independent of the average fitness of its individuals. As the transition proceeds, group fitness gradually becomes "decoupled" from individual fitness (Michod and Nedelcu, 2003), until it is no longer closely related to the average individual fitness. Once group fitness is decoupled, the transition has been achieved, and new complexes have been created that assume an existence of their own.

A New MLS Model

The concept of multilevel selection is very simple: levels are like "Russian matryoshka dolls" (Wilson and Wilson, 2008) nested one within another; selection simultaneously operates on every level and favors different types of adaptations. Many models have been proposed based on this concept (see Wu and Banzhaf (2011) for examples). However, their main focus is to investigate under which conditions the evolution of cooperation will occur or what mechanisms could promote the evolution of cooperation. Furthermore, these models take the hierarchical structure in a MLS for granted; that is, they treat the hierarchical structure as given. Biological hierarchies, on the other hand, have developed gradually; a good example is the evolution of multicellular organisms: it did not exist at the beginning of life. We therefore need to consider other MLS models to explain evolutionary transitions: how simpler entities form complexes and how complexes emerge as discrete units with traits different from their constituents.

This contribution aims at introducing such a new multilevel selection model for evolutionary transitions. The investigation uses the division of labor as an example. Division of labor is a group trait resulting from evolutionary transitions, where low level independent entities with specialized skills cooperate to increase the reproductive success of high level complexes. Examples include the separation of germ and soma cells in simple multicellular organisms, appearance of multiple cell types and organs in more complex organisms, and emergence of casts in eusocial insects (Gavrilets, 2010).

We adopt the extended N-player Prisoner's Dilemma (NPD) game to study the division of labor. The NPD game (Sober and Wilson, 1999) is the classical setting for addressing the evolution of cooperation. Once cooperation is reached, all players possess the same cooperative trait, which is also the only trait required for cooperation. Even if such cooperation breaks down by loosing some individuals, the rest are still capable of cooperating with others. Evidently, the game does not serve the need for investigating the division of labor unless extensions are made. We first change the NPD game by attaching a new trait called "skill" to each player; then we redefine the goal of the NPD game: find N players who not only are willing to cooperate but also possess all required skills.

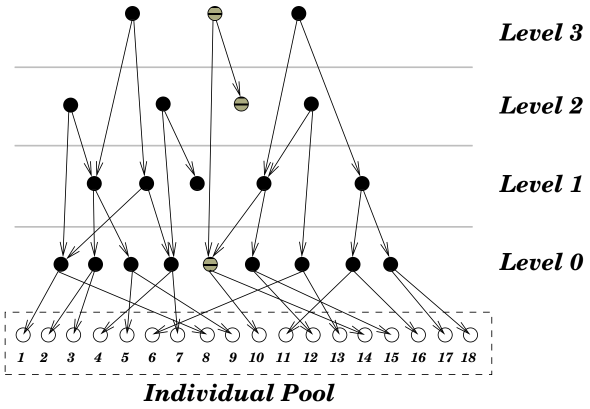

Figure 1: A general framework of the new MLS model

The general framework of our model is illustrated in Fig. 1. This model accommodates two types of entities: individuals (white circles) and groups (black circles). The initial population contains individuals and groups on level 0, which are composed of two randomly selected individuals. The genome of individuals carries two genes. One gene has two variants (alleles); one allele codes for cooperators, the other for defectors. When the former trait is expressed, the individual is said to be a cooperator; otherwise, it is a defector. The other gene encodes the skill possessed. An individual's fitness is determined by the following equations, depending on whether it is a cooperator (C) or a defector (D):

$$f_{\mathrm{C}_i}(x) = \mathit{base} + w\left(\frac{b(n_i q_i - 1)}{n_i - 1} - c\right), \quad (0 \le i < m) \qquad (1)$$

$$f_{\mathrm{D}_i}(x) = \mathit{base} + w\frac{b n_i q_i}{n_i - 1}, \quad (0 \le i < m) \qquad (2)$$

where m is the number of groups in the population; base the base fitness of cooperators and defectors; q_i the fraction of cooperators in group i; n_i the size of group i; b and c are the benefit and cost caused by the altruistic act, respectively; w is a coefficient. From the above fitness definitions, it becomes clear why the initial population must contain groups on level 0: those groups are the smallest units in which the individual fitness can be evaluated. This fitness definition also implies that cooperation is not supported at the individual level, as cooperators always have lower fitness than defectors. Because individuals are unaware of what skills are needed without higher level entities being formed, the skill trait has no effect on the individual fitness.

Groups in the evolution of cooperation simply pool individuals together; however, groups in our model have their own genotype definition, which is represented by a boolean list. Each position in the list is connected to a unique skill, so that the genotype of a group can keep track of all different skills of its members. When a skill is possessed by at least one cooperator in a group, the corresponding position in the genotype is set to true (we say is activated); when the skill is no longer possessed by any cooperator in that group, we inactivate the position by setting it to false. Again, compared to groups in the evolution of cooperation, groups here require their members to develop different skills, not just to cooperate. As a result, groups exhibit more traits than simply the cooperative trait of individuals. Genetically, groups in our model are ready for evolutionary transitions.

From level 0, an operator called "**cooperation**" starts to build the hierarchical structure level by level. In each generation, it selects two existing groups proportional to fitness to form a new group. For example, as highlighted in Fig. 1, a group on level 0 and a group on level 2 can be made to cooperate in a new group on level 3. After the cooperation operator is applied, the genotype of the new group contains all unique skills from the two parent groups. This operator allows evolution to tinker with varying group memberships in order to find the best combination of individuals and groups at lower levels for a higher level function. It in fact is a genetic operator for selecting and reproducing groups; therefore, heritable traits of groups can pass from parent groups to offspring groups. Other genetic operators, such as crossover and mutation, can also be applied to groups. Because groups should be the objects of evolution, multilevel selection of the MLS2 type is employed here.

Group fitness is defined as follows.

$$g(y) = \frac{\sum_{i=0}^{n} f_{\mathrm{idv}}(x_i)}{n} \times \frac{active_{\mathrm{geno}}(y)}{length_{\mathrm{geno}}(y)} \qquad (3)$$

It measures the performance of a group in two respects: (i) the average individual fitness of its n members and (ii) the percentage of activated skills in the genotype. The intention behind this fitness definition is straightforward; the first part encourages the appearance of cooperators, as cooperators improve the overall individual fitness, and the second part rewards groups in which cooperators possess as many different skills as possible. Obviously, this group fitness is not defined as the average individual fitness, but it can be either proportional to average individual fitness, or completely "decoupled" from individual fitness, depending on the influence of the second term of the fitness function. According to Okasha (2005), the former indicates the transition from MLS1 to MLS2, and the latter indicates the groups have fully emerged as discrete units. Both encourage evolution to reach transitions.

Individuals also evolve. To do so, a group is first selected proportional to fitness; an individual is then selected from this group as a parent. For simplicity, asexual reproduction is considered here. Obviously, even though the survival of individuals is now associated with the performance of their group, individuals at this stage are the objects of evolution. Groups provide context for individual fitness evaluation and selection. Hence, multilevel selection of type MLS1 is applied here.

The specific computational implementation of the framework is shown in Algorithm 1. It begins with initialization. N individuals, r percent of which are cooperators, are randomly created and exclusively paired into groups at level 0. Groups at level 0 have their fitness evaluated right away.

In each generation, only one group is created by the cooperation operator, which selects two groups proportional to fitness to create a new group. The consequence of cooperation is the increase of group complexity or the appearance of new levels in the hierarchical structure. To prevent levels from ceaselessly growing, we assign every individual a unique number as its ID; no individuals with the same ID can appear within the same group. After fitness evaluation, the new group is added to the population P. If at that point the maximum number of groups, say N', is reached, another group has to be removed from the population selected inversely proportional to fitness. To highlight the effect of the cooperation operator, crossover and mutation on groups are currently not included.

Algorithm 1: Computational Implementation of the New Multilevel Selection Model

```
P ← Initialize_Population($N$, $r$);
Evaluate_Individual_Fitness($P$);
Evaluate_Group_Fitness($P$);
while population does not converge or max generation is not reached do
    $gp$ ← Conduct_Cooperation($P$);
    Evaluate_Individual_Fitness($gp$);
    Evaluate_Group_Fitness($gp$);
    Add_a_Group_to_Population($gp$, $P$);
    if Population_Size($P$) $> N'$ then
        Remove_a_Group();
    end
    for $i$ ← 0 to $n$ do
        $idv$ ← Reproduce_an_Individual($P$);
        Replace_an_Individual($idv$, $P$);
        Update_Changes($idv$, $P$);
    end
end
```

We also asexually reproduce n individuals every generation. Individuals are selected proportional to fitness from another selected group, instead of from the pool of individuals. The offspring inherits its parent's genome, and further replaces the genome of a less fit individual in the individual pool. The absolute fitness of individuals in the pool is determined by the average fitness of its copies (i.e. individuals with same ID) in all groups. Individuals from the pool are allowed to participate in composing more than one group, so they may have multiple copies in different groups. Depending on group composition, they have different fitness within groups. So the simplest way to determine their absolute fitness is to average the fitness of all copies.

After an individual in the pool is replaced, the change needs to be implemented in all groups that contains the copy of the replaced individual. The group fitness and individual fitness of affected groups need to be updated, accordingly.

We repeat the process until a termination condition has been reached or the population converges.

In summary, this new model distinguishes itself from other multilevel selection models in two ways. First, it integrates two types of multilevel selection, both of which are believed to be relevant to the evolutionary transitions (Okasha, 2005). Individual evolution with the help of group selection is analogous to multilevel selection type 1 (MLS1). It propagates cooperators in the population, which is a prerequisite of evolutionary transitions. Group evolution is then analogous to multilevel selection type 2 (MLS2). The selection pressure on group levels forces groups to evolve adaptations for regulating conflicts among their members. The adaptations indicate that groups emerge as discrete entities with heritable traits. Second, instead of taking the hierarchical structure resulting from evolutionary transitions for granted, our model introduces a "cooperation" operator to create higher level complexes out of simpler ones.

Experiments

In the experiments, we closely examine the transition by our multilevel selection model to the division of labor from a population of independent individuals. First, we examine the ability of our model to evolve groups fulfilling various numbers of skills, when all skills receive the same reward. Second, we examine the dynamics within the model and the responses of individuals when different skills are given different rewards.

Experimental Setup

The experiments are conducted on the extended NPD game with a population of 200 individuals and a maximum of 50 groups on level 1 and above. The initial fraction of cooperators in the population is 0.5; half of the individuals play cooperators in the game, while the other half are defectors. Eq. 1 and Eq. 2 are used to calculate the fitness of cooperators and defectors within a group, respectively. The base fitness $base$ is set to 10, benefit b to 5, cost c to 1, and coefficient w to 1 in these two equations[2]. Group fitness is calculated according to Eq. 3. Group size is a self-adaptive parameter affected by the cooperation operator.

Because the purpose of these experiments is to study the division of labor, our investigation will focus on the effects of two parameters: the number of desired skills and the rewards associated with each skill. For each parameter setting, we ran the model 20 times, each with 5000 generations. We measure the performance of the model by the probability of fixation to cooperators $P_{fixation}$ and the number of activated skills $S_{activated}$. $P_{fixation}$ is computed as the ratio of the number of runs where population converges to cooperators over 20 runs. We also collect the convergence speed $S_{converge}$ in each run, which is the number of generations after which group fitness stops to change.

Varying Skills

The first experiment is given 5 different skills. At initialization, individuals independently choose to be a cooperator or a defector. In addition, they need to randomly pick a skill from 5 skills, $\{1, 2, 3, 4, 5\}$. An individual with an attached skill will perform a specific task. The best performing group should contain only cooperators and should have all 5 skills presented in its genotype. We then gradually increase the number of desired skills to 10, 15 and 20. For each setting, we run the algorithm 20 times. The results are collected in Table 1. The probability of fixation $P_{fixation}$ with a value of 1 is obtained under all settings, which indicates that defectors, despite a relatively high individual fitness, are eliminated from the population, whereas cooperators dominate the population eventually. MLS1 is the explanation for this result. More importantly, the best performing group for each setting develops all required skills through evolution. This demonstrates that MLS2 is at work. It is not surprising to see the larger the number of desired skills, the slower the population was to reach the equilibrium on group fitness. This is simply a reflection of the problem becoming harder when the number of desired skills is raised.

To get a better idea of how the division of labor develops through evolution, we select a typical run for each of $\{5, 10, 15, 20\}$ roles for further analysis. Figure 2 depicts the maximum and average number of unique skills of all groups over 500 generations. Starting from at most 2 skills, the best performing groups gradually evolve to perform more and more different skills until the number of desired skills is reached (see Fig. 2a). This growth is due to the guidance provided by the group fitness. Take the run for 20 desired skills for example. We collect the following information from this run: group fitness, the number of activated roles, and the percentage of cooperators in the best performing group, as well as the percentage of cooperators in the population; that is plotted in Fig. 3.

Group fitness (refer to Eq. 3) is determined by the average individual fitness and the percentage of activated skills. We plot the percentage of cooperators, instead of the average individual fitness, in the best group because of two reasons; we can easily extrapolate the average individual fitness from this percentage, and it also shows the fixation process in the best group. Figure 3 clearly shows how the percentage of cooperators and the number of activated roles affect the group fitness. Interestingly, we notice that the population converges to cooperators first, and then the best group develops all required skills. The same trend is also observed in other runs with $\{5, 10, 15\}$ skills. This observation indicates that cooperators spread in the population before the evolutionary transition happens, a result confirming the discussion about the relationship between MLS1 and MLS2. Group fitness, in turn, influences the execution of individual evolution and group evolution (i.e. cooperation operator). Since defectors bring no fitness benefit on group levels, they are eliminated from the population by group selection at reproduction; hence the percentage of cooperators in the best group and in the population increases steadily towards 1. As shown in Fig. 2b, the average number of activated skills never comes close to the number of desired skills. This implies that the population maintains groups with various

---

[2] Sensitivity analysis of our model wrt. the initial fraction of cooperators and selection pressure (w), as well as a performance comparison with an improved Traulsen's group selection model (Wu and Banzhaf, 2011) can be found in (Wu, 2011). These experiments confirm that our model promotes cooperation over a wider range of parameter settings.

Settings	$P_{fixation}$	$S_{activated}$	$S_{converge}$
$role = 5$	1	5	96.3
$role = 10$	1	10	181.55
$role = 15$	1	15	247.60
$role = 20$	1	20	301.25

Table 1: The performance of our multilevel selection model when individuals play various skills.

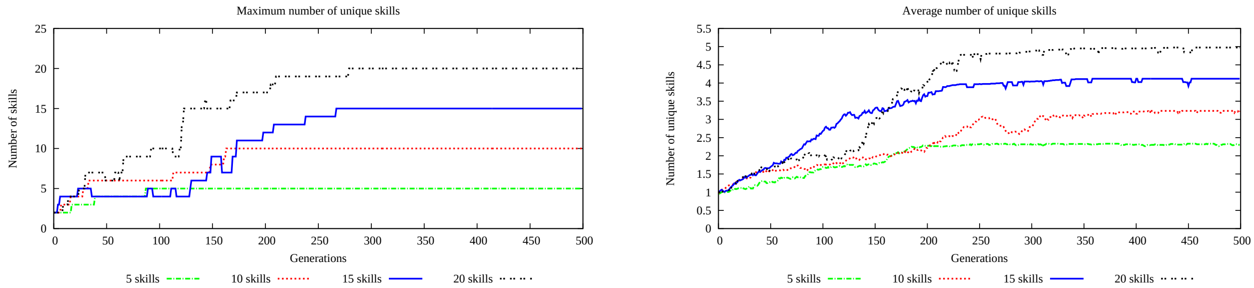

(a) Maximum number of unique skills

(b) Average number of unique skills

Figure 2: The changes of the maximum and average number of unique skills in a typical run.

skills. They are potential building blocks, out of which the cooperation operator is able to test different combinations of existing groups, and gradually hones in on optimal groups with all required skills.

In summary, our model is able to successfully evolve groups with all desired skills for the extended NPD game; or we can say that our model is able to evolve groups to engage in the division of labor between equally rewarded skills.

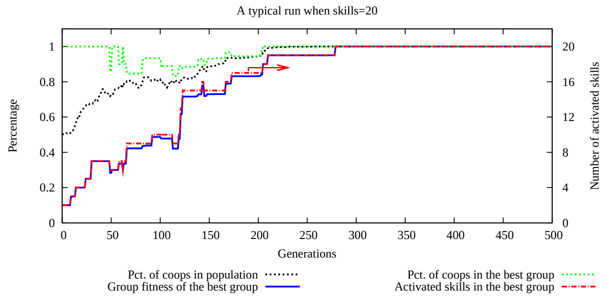

Figure 3: The changes of group fitness, percentage of cooperators and activated roles when 20 skills are set.

Varying Rewards

We continue the exploration of whether or not our model can evolve the division of labor, but this time skills are unequally rewarded. The different rewards put extra pressure on accomplishing the task, as it attracts individuals to specialize on the most rewarding skills while avoiding the less rewarding skills.

To distinguish skills with different rewards, we refer to the "leader/follower" situation described by Goldsby et al. (2009). Individuals who have skill 1 are appointed as leader of that group, while individuals performing other skills are followers. Leaders receive different reward than followers, but followers, no matter what specific skills they have, receive no other rewards. A coefficient, α, is used to control how much reward a leader can receive. Coefficient α basically is a multiplicative of the individual fitness; the individual fitness of a leader is calculated as the product of α and the individual fitness obtained by Eq. 1 or Eq. 2.

We vary the value of α in the range of $\{0.5, 2, 4, 8, 64\}$ on each of $\{5, 10, 15, 20\}$ roles, and run the model on each setting 20 times. The performance is summarized in Table 2. Clearly for each setting the population converges to cooperators as a result of MLS1, and the best performing group is composed of cooperative individuals with all required skills as a result of MLS2.

Because the group fitness can hardly converge in this experiment, the convergence speed $S_{converge}$ is judged by the stabilization of $P_{fixation}$ and $S_{activated}$. Fig. 4 displays a typical run when the number of desired skills is set to 5 and coefficient α is set to 8. Although the percentage of cooperators in the population and the number of activated skills in the best group converge quickly (around generation 350), group fitness and the percentage of leaders in the best group never stop increasing. After generation 350, the percentage of leaders is the only factor that changes the group fitness. Leaders in this case receive much higher rewards than followers, and maximizing this percentage at the same time maximizes the group fitness. Therefore, both values

Settings		$P_{fixation}$	$S_{activated}$	$S_{converge}$
role=5	$\alpha = 0.5$	1	5	90.45
	$\alpha = 2$	1	5	145.35
	$\alpha = 4$	1	5	193.00
	$\alpha = 8$	1	5	238.10
	$\alpha = 64$	1	5	330.00
role=10	$\alpha = 0.5$	1	10	152.2
	$\alpha = 2$	1	10	232.40
	$\alpha = 4$	1	10	379.05
	$\alpha = 8$	1	10	488.00
	$\alpha = 64$	1	10	607.75
role=15	$\alpha = 0.5$	1	15	196.60
	$\alpha = 2$	1	15	313.80
	$\alpha = 4$	1	15	531.50
	$\alpha = 8$	1	15	696.55
	$\alpha = 64$	1	15	950.55
role=20	$\alpha = 0.5$	1	20	314.80
	$\alpha = 2$	1	20	407.35
	$\alpha = 4$	1	20	586.85
	$\alpha = 8$	1	20	902.35
	$\alpha = 64$	1	20	1394.75

Table 2: The performance of runs when leaders are assigned with various rewards.

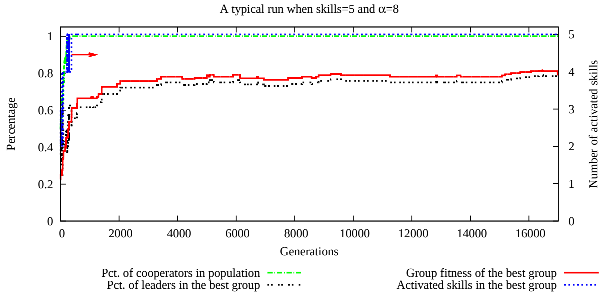

Figure 4: A typical run when skills=5 and α=8.

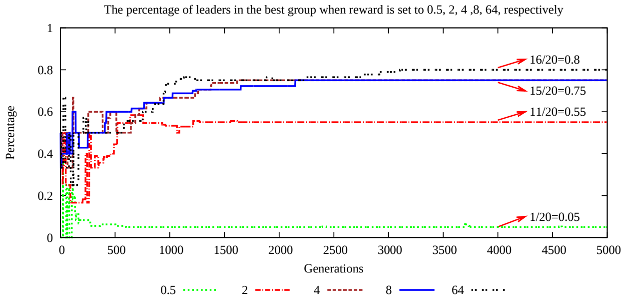

Figure 5: The percentage of leaders in the best group when α is set to 0.5, 2, 4 ,8, 64, respectively.

are constantly improving. Because there is no upper bound on group size, the cooperation operator keeps creating larger groups with more leaders; therefore an equilibrium distribution of different roles can hardly be reached.

To facilitate the investigation on how different rewards affect the division of labor, we restrict the maximum group size to 20. We plot in Fig. 5 the percentage of leaders in the best performing group collected from a typical run with 5 desired skills when α is set to each of $\{0.5, 2, 4, 8, 64\}$. When α is set to 0.5, 5% of 20 individuals, which is only 1 individual, play the role as a leader, while when α is set to 2, 55% of the group, that is 11 individuals, choose to be a leader; similarly, 15 out 20 individuals (75%) become the leader when α is 4 or 8, and 16 leaders (80%) when α is 64.

When α is less than 1, leaders are in fact receiving a penalty, not a reward. Very naturally, individuals avoid becoming a leader, but because of the selection pressure on the group level, the role of a leader must be present in a group. Therefore, the best group ends up with only 1 leader, which maximizes the group fitness. By contrast, when α is greater than 1, individuals strive to be leaders because of the positive reward. An α value of 64 shows another extreme distribution of different roles. Driven by such a significant reward, the best group only has 4 individuals as followers, each for the rest 4 skills, while all other individuals play the role as a leader. The higher the reward, the greater the number of leaders in a group, and the slower the population converges (see $S_{converge}$ column in Table 2).

The experiment perfectly shows the adaptability of our model in response to changes in group selection pressure, and the importance of selection pressure on group levels in developing division of labor. Selection pressure eliminates defectors from a population, adjusts the distribution of roles according to the received reward or penalty, and forces all skills to be present even though some of them have lower fitness than others.

Conclusion

In this paper, we considered a new multilevel selection model to investigate evolutionary transitions. This model introduces a genetic operator called "cooperation" to create higher level complexes out of simpler ones of lower levels. Different types of selection, MLS1 and MLS2, are integrated in the model to determine whether or not the complexes are able to transit to discrete units with their own heritable traits. We test the transition ability of the new model on an extended N-player Prisoner's Dilemma game for achieving the division of labor from a population of independent individuals. The experiments confirm that our model is able to evolve groups fulfilling various numbers of skills whether skills are equally rewarded or not. The experiments also demonstrate that multilevel selection, both MLS1 and MSL2, are necessary for transitions to occur. MLS1 propagates cooperators in a population. Only when participating individuals are willing to cooperate, will evolutionary transitions occur. MLS2 forces complexes to evolve adaptations for regulating conflicts among their members. The adaptations are guided by group fitness, which in our model is decoupled from individual fitness to promote the appearance of new group traits. In future work, we seek to adapt this model for evolutionary computation to solve problems where transitions are needed.

Acknowledgments

W.B. would like to acknowledge support from NSERC Discovery Grants, under RGPIN 283304-07. The authors would also like to thank anonymous reviewers for their helpful comments and suggestions.

References

Borrello, M. E. (2005). The rise, fall and resurrection of group selection. *Endeavour*, 29(1):43–47.

Buss, L. (1987). *The Evolution of Individuality*. Princeton University Press, New Jersey.

Damuth, J. and Heisler, I. L. (1988). Alternative formulations of multilevel selection. *Biology and Philosophy*, 3(4):407–430.

Gavrilets, S. (2010). Rapid transition towards the division of labor via evolution of developmental plasticity. *PLoS Comput Biol*, 6(6):1–10.

Goldsby, H. J., Knoester, D. B., Clune, J., McKinley, P. K., and Ofria, C. (2009). The evolution of division of labor. In *10th European Conference on Artificial Life, Budapest, Hungary, September 2009.*

Michod, R. (1999). *Darwinian Dynamics, Evolutionary Transitions in Fitness and Individuality*. Princeton University Press.

Michod, R. E. and Nedelcu, A. M. (2003). On the reorganization of fitness during evolutionary transitions in individuality. *Integrative and Comparative Biology*, 43(1):64–73.

Okasha, S. (2001). Why won't the group selection controversy go away? *British Journal for the Philosophy of Science*, 52(2001):25–50.

Okasha, S. (2005). Multi-level selection and the major transitions in evolution. *Philosophy of Science*, 72:1013–1028.

Okasha, S. (2008). Units and levels of selection. In Sarkar, S. and Plutynski, A., editors, *A Companion to the Philosophy of Biology*, chapter 8, pages 138–156. Oxford: Blackwell.

Smith, J. M. and Szathmáry, E. (1995). *The Major Transitions in Evolution*. Oxford University Press.

Sober, E. and Wilson, D. S. (1999). *Unto Others: The Evolution and Psychology of Unselfish Behavior.* Harvard University Press.

Traulsen, A. and Nowak, M. A. (2006). Evolution of cooperation by multilevel selection. *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*, 103(29):10952–10955.

Wilson, D. S. (1975). A theory of group selection. *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*, 72(1):143–146.

Wilson, D. S. (1983). The group selection controversy: History and current status. *Annu. Rev. Ecol. Syst.*, 14:159–187.

Wilson, D. S. and Wilson, E. O. (2007). Survival of the selfless. *New Scientist*, 3:42–46.

Wilson, D. S. and Wilson, E. O. (2008). Evolution for the good of the group. *American Scientist*, 96:380–389.

Wu, S. X. (2011). *A Route to the Evolution of Cooperation: Investigations of Multilevel Selection in Evolutionary Computation*. PhD thesis, Department of Computer Science, Memorial University.

Wu, S. X. and Banzhaf, W. (2011). Investigations of Wilson's and Traulsen's group selection models in evolutionary computation. In Kampis, G., Karsai, I., and Szathmáry, E., editors, *Proceedings of the 10th European Conference on Artificial Life, Budapest, Hungary, Sept. 13-16, 2009, Part II*, volume 5778 of *Lecture Notes in Computer Science*, pages 1–9. Springer-Verlag Berlin Heidelberg.