

Mating Restriction and Niching Pressure: Results from Agents and Implications for General EC

R.E. Smith and Claudio Bonacina

The Intelligent Computer Systems Centre, Faculty of Computing
Engineering, and Mathematical Sciences, The University of The West of England, Bristol, UK
{robert.smith,c2-bonacina}@uwe.ac.uk
<http://www.cems.uwe.ac.uk/icsc>

Abstract. This paper presents results and observations from the authors' continuing explorations of EC systems where population members act as autonomous agents that conduct their own, independent evaluation of (and reproduction with) other agents. In particular, we consider diversity preservation in one such agent-based EC system, applied to the multi-peak functions often used to illustrate and evaluate the effects of fitness-sharing-like schemes in GAs. We show how (somewhat surprisingly) *mating restriction alone* yields stable niching in agent-based EC. This leads to a consideration of niching as a generalized phenomenon, and the introduction of *niching pressure* as a concept that parallels *selective pressure*, and which can yield insight. The utility of the niching pressure concept for general EC is explored, and directions for further research are discussed.

1 Introduction

The maintenance of diversity in an EC population is often the key to success in a given task. Diversity can be useful to prevent premature convergence. It is essential in multimodal function optimization, in multi-objective optimization, in dynamic function optimization, and in co-evolutionary machine learning (e.g., learning classifier systems). In many such systems, diversity is maintained through some sort of “niching” operator; an operator that encourages the population to break into separate sub-populations that occupy distinct “niches” of the search space. Mahfoud [8] defines niching methods as “... techniques that promote the formation and maintenance of stable subpopulations in the GA. Niching methods can be applied to the formation and maintenance of interim sub solutions on the way to a single, final solution. They are traditionally viewed, however, in the context of forming and maintaining multiple, final solutions”. Despite the often-essential need for diversity preservation through niching, there is little understanding of the *dynamics* of niching operators.

The authors have recently investigated agent-based EC systems, where each population member acts as an autonomous agent that exchanges information with other such agents, evaluates them as potential mates, and conducts reproduction with them [7,13,14,15]. As has been previously pointed out by the authors, the dynamics of such systems can be non-trivially different from those of centralized EC, and may yield

new insights into agent systems that are not specifically exploiting EC. Also, as will be shown in this paper, they can provide insight into EC in general.

In the following sections, we will discuss the problem of implementing diversity preservation (niching) mechanisms in an agent-based EC system, the testing of such mechanisms on the simple test problems extensively investigated by Deb and Goldberg [3], and results. The results will show that, although mating restrictions alone cannot yield stable niches in a centralized EC system [8], it can in an agent-based EC system. However, this result depends on a subtle balance of effects. Therefore, we introduce the notion of *niching pressure* as a phenomenon that parallels *selective pressure* in an EC system that has niching operators. The broader uses of the niching pressure concept are outlined, and future research directions are discussed.

2 Niching Operators in an Agent-Based EC System

One of the goals of the authors' current research program is to examine systems where EC operates as a means of interaction of individual, autonomous agents, with no centralized EC *per se*. Thus, in the system we consider here (which we will call the ECoMAS model, for **evolutionary computation multi-agent system**) we want to avoid any form of centralized operations. This complicates the implementation of some commonly used niching operators. We will briefly review common niching operators to clarify these issues.

One of the most successful methods to maintain diversity in EC is fitness sharing [6]. Given that this scheme is well known, it will only be briefly discussed here. Fitness sharing is a fitness-scaling scheme that is applied just before parent selection. Assuming application to a multimodal maximization problem, such schemes can be described as means of forcing similar individuals to share their payoff or fitness. This has the effect of limiting the number of individuals in any area of the fitness landscape, based on the relative height of the fitness peaks in that area. Theoretically, fitness sharing should distribute the number of individuals in various areas proportionally to the height of peaks in those areas. With a limited population size, only the "highest" regions will be covered. Fitness sharing works well with fitness proportional selection schemes. Tournament selection decreases the stability of the fitness sharing algorithm, but Oei et al. [12] proposed a solution to combine binary tournament selection and fitness sharing.

The adoption of fitness sharing in the ECoMAS model is problematic.

- Having a centralized perception of the fitness for all individuals is against the decentralized philosophy of our ECoMAS model. We could in principle design a distributed version of fitness sharing or apply a scheme similar to the one described by Yin and Gemany [17]. However, the asynchronous nature of the system would make the implementation of such schemes difficult.
- We do not intend to restrict our system to use non-negative fitness measures, and positive fitness values are an implicit assumption in most fitness sharing schemes. We could rescale fitness values to positive values, but a local imple-

mentation of fitness sharing (necessary in our scheme) would lead to different scaling in different subsets of individuals.

- In our system, individuals locally select their mating partners based on the limited number of *plumages* they receive (in an asynchronous fashion) from other individuals. This is most similar to a tournament selection in regular GAs. We could use the scheme suggested by Oei et al. [12] for combining tournament selection and fitness sharing, but this could also interfere with schemes we might adopt to solve the problems highlighted in the previous bullet points.

Another well-know method to maintain diversity is *crowding* [4]. Crowding works by assuring that new individuals replace similar individuals in the population. As in fitness sharing, similarity is defined by some distance measure between individuals. Unlike fitness sharing, crowding does not allocate solutions proportionally to the height of the peaks. The original crowding scheme proposed by De Jong [4] suffers from replacement errors, which prevent it from stably maintaining solutions close to desired peaks.

Mahfoud [9],[8] modified crowding in order to minimize replacement errors. In the new scheme, called *deterministic crowding*, the population is divided into $N/2$ *random* couples of parents. Two offspring are then generated from each couple. Two parent-offspring pairs are created, minimizing the sum of parent-to-offspring distance. For each pair, the most fit of the two individuals is copied into the new population.

Deterministic crowding is certainly more suitable to adoption in ECoMAS models than fitness sharing. Nevertheless, the following problems are still present:

- In order to implement deterministic crowding, we would have to enforce some level of centralization and synchronization in the system to ensure that each agent at each generation would father only one offspring. This would require centralized information management. We would also have to bind the two agents behaving as parents in our system, because (in principle) we could be willing to replace both of them with their offspring. In order to do so, an undesired level of synchronization would be introduced.
- Co-evolutionary scenarios are one of the aims of our agent-based EC research, and it is not clear what the effects of deterministic crowding would be in such settings.

Deb and Goldberg [3] introduced a mating restriction mechanism as an addition to fitness sharing, in order to prevent recombination between individuals in different niches that may result in low fitness offspring. Deb and Goldberg used a distance measure in the decoded parameter space (*phenotypic mating restriction*) and a threshold σ_{Mating} . Deb and Goldberg then set $\sigma_{\text{Mating}} = \sigma_{\text{Share}}$ for simplicity. Only individuals that are less than σ_{Mating} apart from each other are, in principle, allowed to mate. If no such individual is found, a mating partner is chosen at random. Deb and Goldberg found that the introduction of phenotypic mating restrictions enhanced the on-line performance of the fitness sharing algorithm.

Mahfoud [8] has proven that the mating restrictions introduced by Deb and Goldberg cannot maintain diversity in a centralized EC system on their own. However, this is not the case in agent-based EC, as we will later see. Moreover, mating restrictions have a straightforward, entirely local implementation in our agent-based scheme.

3 Preliminary Results on Mating Restriction in Agent-Based EC

This section describes a very basic test of our ECoMAS model with mating restrictions. In order to conduct a basic test of mating restrictions in agent-based EC, we use some of the multimodal functions described in [3]. Specifically, we consider the following two multimodal functions:

$$F1: f_1(x) = \sin^6(5\pi x), \quad F2: f_2(x) = e^{-2\ln 2 \left(\frac{x-0.1}{0.8} \right)^2} \sin^6(5\pi x)$$

F1 has five peaks of equal height and width, evenly spread in the range $0 \leq x \leq 1$. *F2* has five peaks of decreasing height and equal width, evenly spread in the range $0 \leq x \leq 1$.

In the ECoMAS system applied to these simple test functions, individuals send out plumages and collect plumages. A plumage is composed of features of the sending agent. This is typically its real-valued decoding (the “phenotype”). The fitness function value and genotype are also included for convenience, but we do not, in general, assume that a centralized fitness value is available, or that the genotype can be inferred from the plumage for purposes of reproduction, so that the decentralized intent of the ECoMAS model is not compromised. Each agent sends out *number_of_receivers* plumages initially, and each time it receives a plumage. Each agent gathers a list of plumages, which represents its list of possible partners. As soon as the list contains *max_partners* agents, the individual can behave as a “mother” and generate an offspring. An agent acting in the mother role sorts the list according to its own (fitness and niching) criteria. The agent tries to get the “best” individual in her list to father her offspring, by requesting his genetic material. If the father is no longer available in the system, the mother tries to get to the second best father in its list, and so on. If none of the individuals in the list are available, the mother produces the offspring asexually (via mutation only). If a father is successfully found, the offspring is the result of the recombination of the genotype of the mother and the father. After an agent produces *max_children* offspring in the mother role, it dies (eliminates itself from the system).

In the mating restriction scheme adopted here, each agent filters the individuals that it inserts into its list of partners. If an agent receives a plumage from an agent that is farther than a local threshold (*mating restriction threshold*) away (according to some distance metric), this plumage is not added to the partner list, and is simply discarded. We consider two distance metrics:

- Euclidian distance, based on the real-valued decoding of the binary string representing each individual. When using this metric we will call the mating restriction scheme *Euclidean mating restrictions*.

- Hamming distance, based directly on the binary genotype of each agent. When using this metric we will call the mating restriction scheme *Hamming mating restrictions*.

The terms *phenotypic* and *genotype* mating restriction (respectively) are used for these sorts of schemes in the literature. However, the authors feel that these terms aren't particularly accurate, given that the distinction in most cases, and particularly in our system, is only that of a neighborhood structure, not the *developmental* distinction of phenotype from genotype.

In the experiments presented in this section, 100 agents, each with a 30-bit genotype, are run for 200 generations, using two-point crossover with probability 0.9, and *number_of_recievers* = 10. Setting *max_children* to one insures the initial number of agents (population size) is constant. A mutation probability of zero is used for comparison of results to Deb and Goldberg’s original experiments.

Results in Fig. 1 show that the mating restriction scheme yields stable niches on each of the peaks. Note that niches are formed in less than 30 generations, and after that time they are absolutely stable, by the nature of the system. This result is somewhat surprising, given that in regular GAs, mating restriction alone is not expected to yield stable niching. This effect is explored further in the following sections.

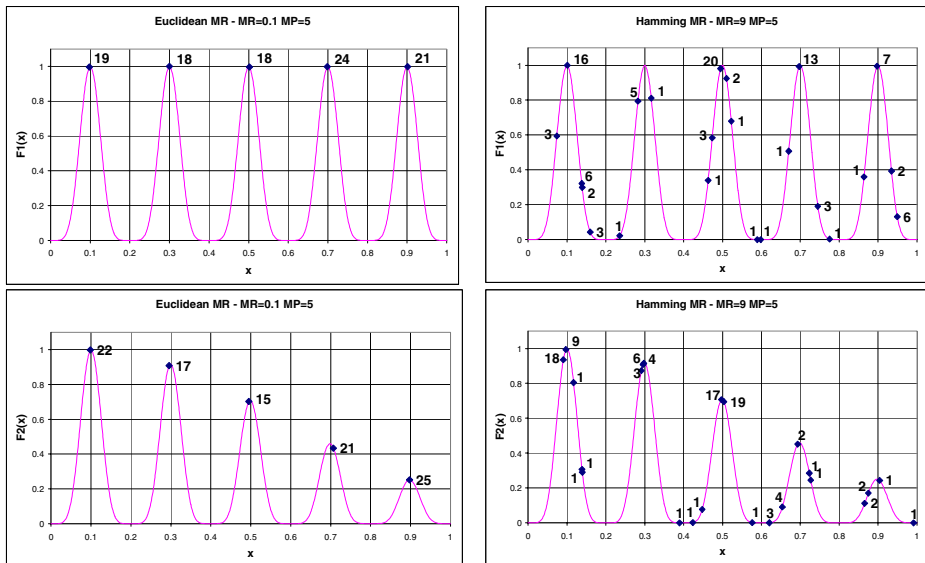


Fig. 1. Results of Mating Restriction in the ECoMAS model. Numbers in these figures are used to indicate the number of individuals on each peak. Settings for *max_partners* and *mating_restriction_threshold* are noted in sub-figure titles

4 Niching Pressure

In the ECoMAS model whose results are shown above, mating restrictions are applied *before* an individual selects its mate. If a mating restriction cannot be fulfilled, the individual in the “mother” role reproduces asexually. Effectively, individuals only interact reproductively (that is, in selection, recombination, and replacement) with individuals within *mating_restriction_threshold* distance of themselves. Over time, individuals that are within this distance from one another *separate* into subpopulations that are effectively isolated from one another. While it might be possible to simulate this effect in a non-agent-based system, it is a natural and desirable side effect of the agent perspective, which may have implications for other EC systems.

The behavior is in strict line with Mahfoud’s [8] comments that “niching methods alter the selection algorithm to provide selective pressure within, but not across regions of the search space. The selective pressure within individual regions can be substantial, and still preserve niching properties.”

Clearly (like with most aspects of EC) there is a subtle balance involved here. If the population divides into isolated subpopulations too quickly, selective pressure can only find peaks in a limited number of areas of the search space (each covered by a subpopulation), and some peaks may not be found. However, if the population does not divide quickly enough, selective pressure (or drift) will drive the population towards too few peaks. In either case, not all peaks will be maintained.

In an attempt to better understand the balances involved, this section will introduce the concept of *niching pressure* as a parallel to selective pressure. Thus, we first review the concept of selective pressure itself.

Selection typically emphasizes better solutions in the population in two different stages of the evolutionary algorithm:

- Selection of parent(s): the operator has to select the parent(s) from which the offspring will originate.
- Selection of survivor(s): after the generation of the offspring, the selection operator has to decide which members of the population are to be replaced by the new ones and which are going to survive.

Selective pressure is informally defined as the emphasis of selection of the best individuals. Selection in GAs typically yields an exponential form of growth in numbers of observed higher-fitness schemata. To provide a characteristic for this exponential growth, Goldberg and Deb [5] and Back [1] related selective pressure to *takeover time*, which Deb [2] defines as “... the speed at which the best solution in the initial population would occupy the complete population by repeated application of the selection operator alone.” Clearly, takeover time, t_t , can vary between ‘1’ and infinity. In Goldberg and Deb [5], selective pressure is inversely proportional to takeover time. The longer the takeover time, the lower the selective pressure; and vice versa. We will assume the selective pressure to vary in the range $[0, \infty)$. If $t_t = 1$ the selective pressure is infinite; if $t_t = \infty$, then the selective pressure is ‘0’.

As noted above, the effect of mating restrictions in the ECoMAS model is for isolated subpopulations to emerge and (in line with Mahfoud’s comments) for selective

pressure to then act only within those subpopulations. We will say that separation into subpopulations has occurred when a final, particular set of isolated subpopulations emerges. By isolated, we mean populations that do not influence selection or recombination in one another. Paraphrasing Deb's description of takeover time, separation time is the speed at which the final set of subpopulations emerges through repeated application of the selection and niching operators alone. In line with Mahfoud's comments, selection can still occur within these subpopulations (that is, particular individual types in a subpopulation may be eliminated) after separation, but we consider separation to be complete when these sets are stable, regardless of their contents.

Given the previous discussion, we will define niching pressure in a fashion similar to selective pressure. Specifically, we will identify niching pressure as inversely proportional to the time until separation into a final set of stable subpopulations (the separation time). Like takeover time and its inverse, selective pressure, separation time and niching pressure are simplified, single-number characterizations of a complex process, which in practical GAs is unlikely to be so simple. However, like selection pressure, niching pressure is a useful concept.

A rise in niching pressure (whose controls we have not yet discussed) should result in a reduction in separation time. However, note that niching pressure and selective pressure are not entirely separate. Raising selective pressure should reduce both takeover time and separation time. Selective pressure drives niching pressure, and niching (separation) ultimately defines the subpopulations in which selective pressure can act.

5 Balancing Selective and Niching Pressure

How is niching pressure controlled? First, consider how selective pressure is controlled in the previously discussed ECoMAS model (we will later generalize to other models). Assuming *max_children*=1, selective pressure is primarily controlled by *max_partners*, the size of the list upon which "mother" agents base their selective decisions, and (less directly) by *number_of_receivers*, the number of plumages an agent sends out each time in the father role. Increasing *max_partners* increases the number of individuals against which any potential father is evaluated, thus reducing the chances of mating for less-fit fathers. Increases in this parameter should lower takeover time, and thus increase the selective pressure. Note that this is similar to the effect of raising the tournament size in tournament selection.

As was discussed earlier, niching pressure is directly effected by controls on selective pressure. We accounted for selective pressure controls in the previous paragraph. However, we also want to account for controls on niching pressure that are separable from controls on selective pressure. Niching pressure is inversely proportional to separation time. Clearly, separation time (but not takeover time within any given subpopulation) is controlled by the *mating restriction threshold*. Thus, this is a control on niching pressure that does not directly affect selective pressure. Generally, the higher the *mating restriction threshold* the longer it will take for separation to occur. In the limiting cases, a threshold that nearly spans the space will result in a very long separation time, and a threshold of zero means the population is fully separated im-

mediately. Thus, we theorize that the *mating restriction threshold* directly controls separation time, and *inversely* controls niching pressure. In the next paragraphs we will test the effects of the controls suggested here.

We consider Euclidean mating restrictions with three levels of the *mating_restriction_threshold* to control niching pressure: 0.05, 0.10, and 0.15. We consider four levels of *max_partners* to control selective pressure. These are 2, 5, 10, and 15. Similar investigations for Hamming mating restrictions have been performed, but those results cannot be presented here, for the sake of brevity, and will be presented in a later document. All results presented are averages over five runs. We adopted the Chi-squared-like measure introduced by Deb and Goldberg [3]. This measures the deviation from a desired distribution. In the case of F1, the desired distribution is equal numbers of individuals at the top of each peak. In the case of F2, the number of individuals on top of each peak in the desired distribution is proportional to the peak's height. This measure is not entirely informative, since the mating restriction scheme does not have the global perspective of fitness sharing, and does not explicitly attempt to “balance” the number of individuals on each peak. Since the mating restriction scheme works differently with regard to peak height (results of which we will later discuss), we evaluate the Chi-square measure against both an equal and a proportional distribution on the peaks for function F2. Like in Deb and Goldberg's study, we consider an individual to be “on” a peak if its fitness is more than 80% of the peak's value.

To more thoroughly consider cases where individuals are represented on every peak, but not necessarily distributed proportionally, we also consider the average number of peaks covered by more than 4 individuals (the expected number in a random population) and more than 10 individuals (the expected number if individuals were evenly placed in the “bins” defined in Deb and Goldberg's Chi-square-like measure).

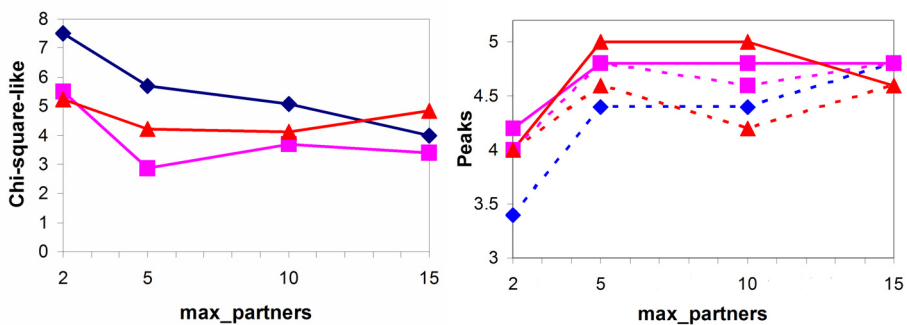


Fig. 2. Results of varying niching and selective pressure, function F1 (equal peaks), Euclidean mating restrictions. *Mating_restriction_threshold* values are 0.05=◆, 0.10=■, 0.15=▲. In the right-hand graph, solid lines are average number of peaks with over 4 individuals; dotted lines are peaks with over 10 individuals. Note coincident solid lines for ◆ and ■ in the right-hand graph

Fig. 2 shows the results for F1 and Euclidean mating restrictions. The graphs show a complex interplay of selective and niching pressure, which can be clarified by other data, which cannot be included here for the sake of brevity. Some explanation will suffice. Although error bars are not included in the graphs for the sake of visual clarity, the following statements are true with 90% confidence levels.

With the highest niching pressure (*mating_restriction_threshold* = 0.05), the population rapidly separates into many small niches. Some of these are off of the function's peaks. Therefore, increase in the selective pressure is necessary to drive individuals up the peaks before separation occurs. Thus, increase in selective pressure results in better Chi-square-like measure (left-hand graph) and in better distribution of individuals on peaks (right-hand graph, dotted line). For slightly lower niching pressure (*mating_restriction_threshold* = 0.1), niches are larger, and separation is slower. This allows for low values of selective pressure to generate better performance than in the previous case. For all but the lowest value of selective pressure, all peaks are represented in almost all runs. If the niching pressure is on its lowest value (*mating_restriction_threshold* = 0.15), coupled with moderate selective pressure, the best performance is obtained. While the “even” distribution indicated by the Chi-square-like measure is slightly worse, all peaks are found with great consistency. However, if the selective pressure is too high, drift occurs before separation, resulting in a worse distribution of individuals on peaks (right-hand graph, solid line). In fact, a few peaks are not even covered by four agents (which we interpret as peak loss). Also, the dotted line goes up, because the individuals that do not cover some peaks end up on over-represented peaks (which we interpret as drift). The drift phenomenon was also observed with Hamming mating restrictions. In that case, the effects were much more pronounced, and were particularly evident even for intermediate values of niching pressure. Drift had a greater effect, due to the more connected neighborhood structure of the Hamming space.

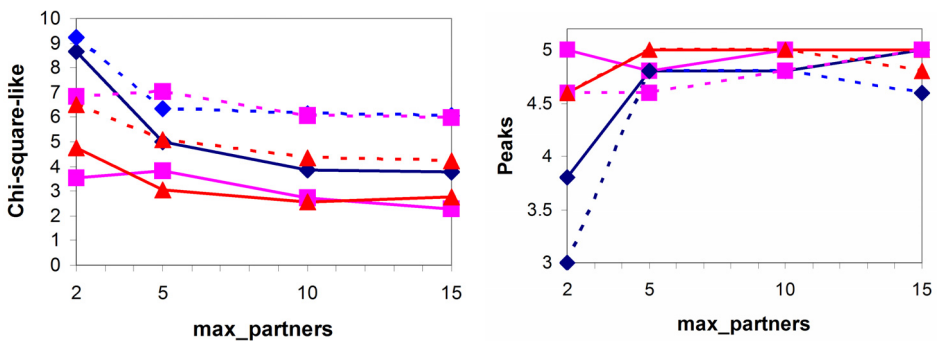


Fig. 3. Results of varying niching and selective pressure, function F2 (decreasing peaks), Euclidean mating restrictions. Symbols are generally as in the previous figure. In the left-hand figure, dotted lines represent deviations from the proportional distribution, and solid lines from the equal distribution

Fig. 3 shows results for F2 and the Euclidean mating restriction. In this case, increasing selective pressure generally enhances results for all levels of niching pressure. However, note that for the highest value of selective pressure ($max_partners=15$) and the highest value of niching pressure ($mating_restriction_threshold = 0.05$), peaks are beginning to be underrepresented (reflected in the right-hand graph) because separation occurs before all peaks represented adequately.

6 Implications for Other Niching Schemes

Given our consistent use of mating restrictions within the ECoMAS model as an illustrative example, it may at first seem that the introduction of niching pressure as a concept is particular to this system. The authors believe this is not the case, and that the concept can be usefully generalized.

We will call the group of individuals over which selection acts the *selection group*. In characterizing what we mean by the selection group, we must consider both roles of selection mentioned previously: selection to act as parents, and selection for replacement. Generally, we observe that *increase in the number of individuals in the selection group increases selective pressure*.

For instance, consider deterministic crowding, where selection to act as parents is random, and selection for replacement is within a parent-offspring couple. Thus, selective pressure could be controlled in a manner similar to tournament selection. One could enlarge the group of parents, or the number of children produced by parents.

We will call the maximum expected range of distances between any pair of individuals in a selection group that group's *span*. We further observe that *increase in the span of the selection group is likely to decrease niching pressure because it will increase separation time between subpopulations*.

In the ECoMAS model, the span of the selection group is directly controlled by the mating restriction threshold, whose effects were illustrated in the previous section. In deterministic crowding, the span is between a parent and the children with which it is paired (based on distance). The span of this group is an artifact of the representation, and the recombination operators. This complicates analysis. It is possible that full separation may never occur in deterministic crowding, since under given operators for a given representation, selection groups from a given subpopulation may always overlap other subpopulations. This may sometimes cause inability of deterministic crowding to maintain stable niches. Also note that controls on selective and niching pressure in deterministic crowding are more difficult to separate than in ECoMAS.

The previously introduced concepts can be generalized to other schemes, for instance, island model GAs [11] and fine-grained GAs [10]. Although these schemes are known to be unable to maintain stable niches, extensions of the concepts reveal the possibility of analogous controls on selective and niching pressures.

In the island model, subpopulations are defined by distinct islands. Selection groups are defined by those subpopulations and migrants from other islands. Note

that in this case, the group for selection of parents could span the same space as the islands (depending on the migration policy), while the group for replacement remains local to each island. Thus, the control on niching pressure is on the span of migrants that may come from other islands. If this was dynamically adjusted towards zero over time, stable niches could be obtained. Note that the number of islands also influences niching pressure, independently of selective pressure. The authors believe that a careful theory of balancing niching pressure and selective pressure should yield insights into appropriate migration tuning policies for desired effects in island models.

In fine-grained GAs, an individual's local neighborhood defines both its subpopulation and its (parental and replacement) selection group. Once again, tuning the selective group size and span in a fashion that balances selective and niching pressure over time should yield desired niching.

Finally, consider the implications of the previously introduced concepts for learning classifier systems [6], particularly XCS [16]. In this system, both subpopulations and (parental) selection groups are defined by the match-and-act process, since this process forms the groups of individuals that participate in the non-panmictic GA. We believe an emerging theory of balance for selective and niching pressure (for which this paper is a start) should clarify niching in XCS.

Fitness sharing (perhaps the most popular niching scheme) [6] is something of an exception. Fitness sharing never separates the population into distinct subpopulations within which selection is isolated, and selective groups always span the population in its centralized scheme. Selection is always possible between all niches. While this sounds like a prescription for drift towards one peak over all others, fitness sharing corrects by globally renormalizing the fitness of peaks based on the number of individuals at each peak. While this is a powerful scheme when centralization is possible, it also explains the difficulty in transferring fitness sharing to more distributed GAs like ECoMAS.

7 Discussion and Final Comments

This paper has introduced results with an agent-based EC system where stable niches emerge through the effects of mating restrictions, and introduced the notion of niching pressure as a parallel to the well-known concept of selective pressure. Niching pressure is defined to be inversely related to separation time, which is in turn defined as the theoretical time until isolated subpopulations emerge. We experimentally considered controls on the balance of these two pressures, and how the concepts revealed can be extended to other systems, through the notion of the size and span of selection groups, in relationship to subpopulations.

Clearly, the results of this paper are preliminary. Thorough analytical and empirical treatments of the concept of niching pressure are now necessary. However, the authors strongly feel that the concepts introduced will provide much greater insight into an important, balancing effect in EC systems with niching.

References

- [1] Back T. (1994) Selective pressure in evolutionary algorithms: A characterisation of selection mechanisms. In *Proceedings 1st IEEE Conference On Evolutionary Computation*. pp. 57–62. IEEE Press.
- [2] Deb, K. (1997) Introduction to selection. In Back T., Fogel D. B. and Michalewicz Z. (eds.) *The Handbook of Evolutionary Computation*. pp. C2.1:1–4. Institute of Physics Publishing and Oxford University Press.
- [3] Deb, K. and Goldberg, D. E. (1989). An investigation of niche and species formation in genetic function optimization. *Proceedings of the Third International Conference on Genetic Algorithms*. p. 42–50. Morgan Kaufmann.
- [4] DeJong, K.A. (1975). *An analysis of the behavior of a class of genetic adaptive systems*. Unpublished Doctoral Dissertation. The University of Michigan.
- [5] Goldberg D.E. and Deb K. (1991). A comparative analysis of selection schemes used in genetic algorithms. In Rawlins G. J. E. (ed.) *Foundations of Genetic Algorithms*. pp. 69–93. Morgan Kaufmann.
- [6] Goldberg, D.E. (1989). *Genetic algorithms in search, optimization, and machine learning*. Addison-Wesley. Reading, MA.
- [7] Kearney, P., Smith, R., Bonacina, C., and Eymann, T. (2000). Integration of Computational Models Inspired by Economics and Genetics. *BT Technology Journal*. v18n4. pp. 150–161.
- [8] Mahfoud, S.W. (1995). *Niching method for genetic algorithms*. Unpublished Doctoral Dissertation. University of Illinois. IlliGAL Technical Report 95001.
- [9] Mahfoud, S.W. (1992). Crowding and preselection revisited. In *The Proceedings of the Second Parallel Problem Solving From Nature Conference*. pp. 27–36. Springer.
- [10] Manderick, M. and Spiessens, P. (1989). In *Proceedings of the Third International Conference on Genetic Algorithms*. pp. 428–422. Morgan Kaufmann.
- [11] Martin, W.N., Lienig, J., and Cohoon, J.P. (1997) Island (migration) models: Evolutionary algorithms based on punctuated equilibria. In Back T., Fogel D. B. and Michalewicz Z. (eds.) *The Handbook of Evolutionary Computation*. pp. C6.3:1–16. Institute of Physics Publishing and Oxford University Press.
- [12] Oei, C. K., Goldberg, D.E., and Chang, S. (1991). Tournament selection, niching, and the preservation of diversity. IlliGAL Technical Report 91011.
- [13] Smith, R. E. and Bonacina, C. (2003). Evolutionary computation as a paradigm for engineering emergent behaviour in multi-agent systems. In Plekhanova, V. (ed.) *Intelligent Agent Software Engineering*. pp. 118–136. Idea Group Publishing.
- [14] Smith, R.E. and Taylor, N. (1998). A framework for evolutionary computation in agent-based systems. In C. Looney and J. Castaing (eds.) *Proceedings of the 1998 International Conference on Intelligent Systems*. ISCA Press. p. 221–224.
- [15] Smith, R.E., Bonacina C., Kearney P. and Merlat W. (2000). Embodiement of Evolutionary Computation in General Agents. *Evolutionary Computation*. 8:4. pp. 475–493.
- [16] Wilson, S.W. (1995). Classifier fitness based on accuracy. *Evolutionary Computation*. 3(2), pp. 149–176.
- [17] Yin X. and Gernay N. (1993) A fast genetic algorithm with sharing scheme using cluster analysis methods in multimodal function optimization. In Albrecht R.F., Reeves C.R. and Steele N.C. (eds.) *Artificial Neural Nets and Genetic Algorithms: Proceedings of an International Conference*. pp. 450–457. Springer.