An Analysis of the Effects of Neighborhood Size and Shape on Local Selection Algorithms

Jayshree Sarma and Kenneth De Jong

Computer Science Department
George Mason University
Fairfax, VA 22030
jsarma, kdejong@gmu.edu

Abstract. The increasing availability of finely-grained parallel architectures has resulted in a variety of evolutionary algorithms (EAs) in which the population is spatially distributed and local selection algorithms operate in parallel on small, overlapping neighborhoods. The effects of design choices regarding the particular type of local selection algorithm as well as the size and shape of the neighborhood are not particularly well understood and are generally tested empirically. In this paper we extend the techniques used to more formally analyze selection methods for sequential EAs and apply them to local neighborhood models, resulting in a much clearer understanding of the effects of neighborhood size and shape.

1 Introduction

Adapting evolutionary algorithms to exploit the power of finely-grained parallel architectures poses a number of interesting design issues. A standard approach is to use spatially structured populations in which local selection algorithms operate in parallel on small, overlapping neighborhoods [4, 7, 8, 9]. The effects of design choices regarding the particular type of local selection algorithm as well as the size and shape of the neighborhoods are not particularly well understood and generally tested empirically.

An EA, whether serial or parallel can be effective only when a proper balance between exploration (via well chosen operators) and exploitation (well chosen selection pressure) is maintained. Having some insights as to how the selection pressure can be varied in local neighborhood EAs will help in designing better parallel EAs.

In De Jong and Sarma 1995 [5] we studied the emergent global selection pressure induced on the entire population by standard local selection algorithms. Using standard small neighborhood sizes and shapes we presented results which indicated that the emergent global selection pressure of a particular local selection algorithm was qualitatively similar to its sequential counterpart, but quantitatively weaker.

In this paper we extend these results by looking more closely at the effects of local neighborhood size and shape. Our approach involves generalizing the techniques used to analyze selection methods for sequential EAs and apply them to local neighborhood models. In particular, we extend growth curve analysis to local neighborhood models resulting in a much clearer quantitative understanding of the effects that neighborhood size and shape have on the emergent global selection pressure.

2 Finely-grained Parallel EAs

There are a variety of finely-grained parallel EAs which have been proposed and studied (see, for example, [2, 4, 7, 10].) For this study we have adopted a fairly standard model to analyze. We assume a two-dimensional toroidal grid as the spatial population structure in which the neighborhood of a particular grid point is defined in terms of the number of steps taken (up, down, left, right) from that grid point. Every grid point has a neighborhood which overlaps with the neighborhoods of nearby grid points, and all neighborhoods are of identical size and shape.

Each grid point contains one individual of the population and, in addition, an evolutionary algorithm is assumed to be running simultaneously on each grid point, continuously selecting parents from the neighborhood of that grid point in order to produce offspring which replace the current individual assigned to that grid point. The overlapping neighborhoods provide an implicit mechanism for migration of genetic material throughout the grid. The amount of overlap is a function of the neighborhood size and shape.

Figure 1 illustrates four of the five neighborhood configurations used in this paper and these are neighborhoods that are typically used in the literature. The shape label L (linear) is assigned to neighborhoods defined as all points reachable in $\leq n$ steps taken in a fixed axial direction (north, south, east, or west) from the central grid point, while C (compact) neighborhoods contain the closest $n-1$ points to the central grid point.

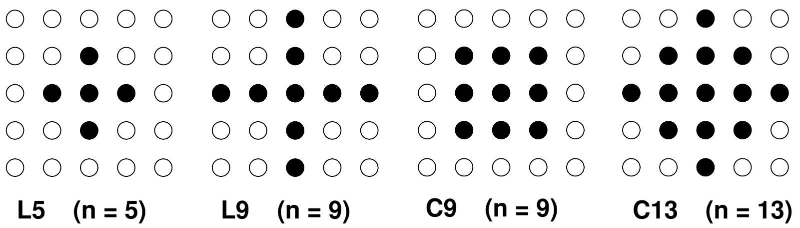

Fig. 1. Neighborhood Patterns

The selection algorithms used on the local neighborhoods are typically the same ones used for sequential EAs. In this paper we focus on just two: fitness proportional and linear ranking selection.

3 Growth Curve Analysis

A standard technique for studying and comparing selection algorithms for serial EAs is to characterize the selection pressure they induce on the individuals in a population in which only reproduction is active (i.e., no mutation, crossover, etc.). Of particular interest is the growth rate of the best individual in the initial population. For the standard selection algorithms, such as fitness proportional and linear ranking, these growth curves are logistic in nature but vary in their growth rates [1, 6]. Since the growth rate of fitness proportional selection is dependent on the fitness ratio, we kept the fitness ratio constant at 2.0 in all our studies in this paper.

Figure 2 illustrates how these growth rates change when we move to spatially structured populations with local neighborhood selection. In this particular case the growth curves exhibited by the best individual in a serial EA with a population size of 1024 using rank and fitness proportional selection are compared with the growth curves obtained from a parallel EA using a 2-dimensional 32×32 toroidal grid and a local neighborhood size of 9.

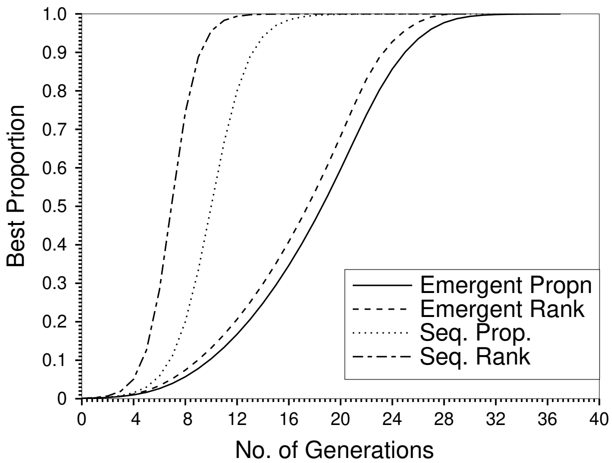

Fig. 2. Sequential (population 1024) and emergent (32×32 grid, neighborhood size 9) growth curves for fitness proportional and linear ranking selection.

Figure 2 is typical of what is consistently observed, namely the spatially structured EAs exhibit the familiar logistic growth curves but with lower growth rates. Note also that the rank order of the selection intensity is preserved: using weaker local selection pressure (in this case, proportional selection) induces a weaker emergent global selection pressure.

The intuitive explanation for this phenomenon is that it is the result of the combined effects of logistic growth occurring in each of the local neighborhoods and the propagation times necessary to spread the best individual globally throughout a spatially structured population. In De Jong and Sarma 1995 [5] we confirmed the first of these two hypotheses showing that the emergent global selection pressure of a local selection algorithm such as fitness proportional or linear ranking selection is qualitatively similar to that produced by fitness proportional or linear ranking in standard serial EAs with global mating pools. In this paper we focus on the second hypothesis involving the effects of propagation times on growth rates.

4 Effect of Neighborhood Size

Increasing the neighborhood size creates a larger overlap and decreases the propagation time. Thus one can see that propagation times are closely related to neighborhood size. Hence increasing local neighborhood sizes while keeping the grid size fixed should result in corresponding increases in selection intensity reflected by higher growth rates of the best individual.

Figures 3 and 4 illustrate typical results obtained on the same 32×32 toroidal grid used in the previous section. For both ranking and fitness proportional selection we analyzed the growth curves for five neighborhoods of different sizes (5, 9, and 13) and shapes (C and L).

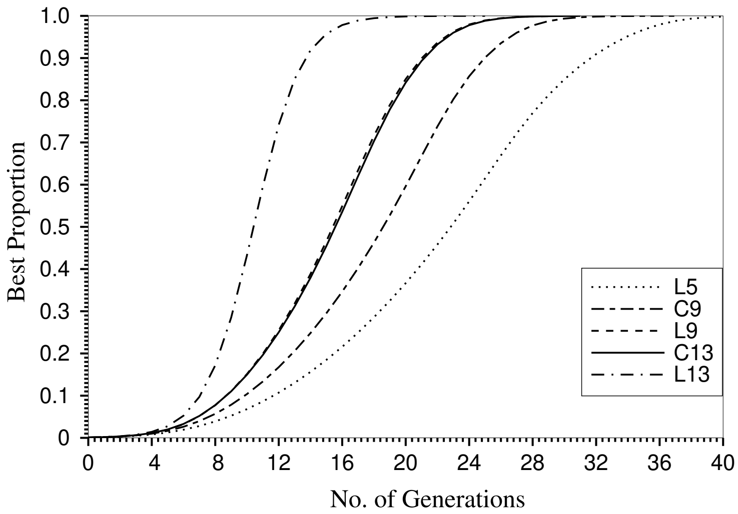

Fig. 3. Emergent growth curves for fitness proportional selection using different neighborhood sizes.

The results are fairly consistent with our expectation that selection intensity increases with increasing neighborhood size. However, there is an interesting "anomaly" in both Figure 3 and Figure 4. The L9 and C13 neighborhoods exhibit nearly identical growth curves even though C13 is almost 50% larger than L9. This suggests that neighborhood shape also plays an important role.

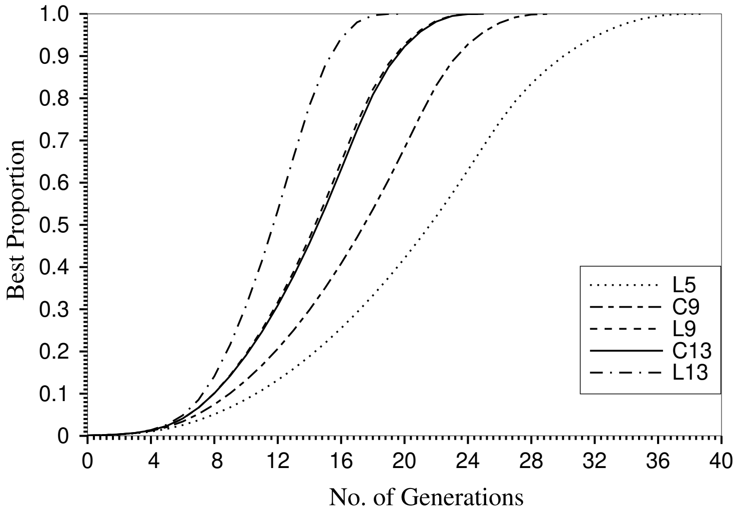

Fig. 4. Emergent growth curves for linear ranking selection using different neighborhood sizes.

5 Effects of Neighborhood Shape

In attempting to explain why the L9 produces a growth rate significantly higher than C9 and nearly identical to C13, the intuitive notion is that what is really important here is the radius of a neighborhood. If we consider the radius of a neighborhood to be defined as the radius of the smallest hypersphere containing the neighborhood, then C9 has a radius of $\sqrt{2}$ while L9 and C13 have radii of 2. This correlates reasonably well with the observed differences in growth rates noted in the previous section.

However, this simple definition of neighborhood radius does not distinguish between L9 and C13 which have identical radii, but have slightly different growth rates resulting from the finer interior details of neighborhood shape. To account for this we adopted an alternative definition for neighborhood radius based on a standard distance measure used in spatial analysis [11]:

$$rad = \sqrt{\frac{\sum(x_i - \bar{x})^2 + \sum(y_i - \bar{y})^2}{n}}$$

which measures the spatial dispersion of a point pattern and can be thought of as the radius of the circle centered on the mean center $(\bar{x}, \bar{y})$ of a neighborhood pattern of n points where

$$\bar{x} = \frac{\sum_{i=1}^{n} x_i}{n}, \qquad \bar{y} = \frac{\sum_{i=1}^{n} y_i}{n}$$

Using this definition, the radii of C9, L9, and C13 are 1.16, 1.49, and 1.47 respectively. A closer look at Figure 3 and Figure 4 indicates that this correlates well with the observed growth curves, with L9 actually exhibiting slightly higher growth rates than C13.

6 Effects of Neighborhood Radius

The results of the previous sections suggest a simple, intuitive model in which selection intensity varies as a function of radius of the local neighborhoods and can be driven arbitrarily close to the selection intensities of the corresponding serial EAs by choosing sufficiently large radii. Figure 5 illustrates this for three different radii on the same 32×32 grid as before using fitness proportional selection. Figure 6 shows the same results for linear ranking selection.

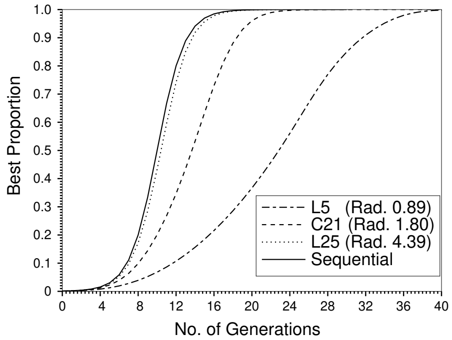

Fig. 5. Emergent growth curves for fitness proportional selection using different radius.

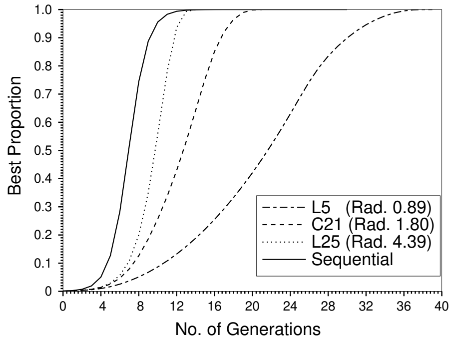

Fig. 6. Emergent growth curves for linear ranking selection using different radius.

If our hypothesis is correct, we should observe similar effects if we keep the neighborhood radius fixed and decrease the grid size. Figures 7 and 8 illustrate that this is precisely what happens when the radius is kept fixed at 0.89 and the grid sizes are changed.

This in turn suggests that the key factor in controlling selection intensity is the ratio of the neighborhood radius to the grid radius. Figures 9 and 10 illustrate that this is precisely the case. Here an L5 neighborhood on a 32×32 grid with a ratio of 0.06815 produces a growth curve nearly identical to a C21 neighborhood on a 64×64 grid with a ratio of 0.06891.

7 A Simple Quantitative Model for Selection Intensity

The observations of the previous section suggest how one might derive a simple quantitative model of selection intensity for these local neighborhood models. Like their sequential counterparts, all of the growth curves generated by the

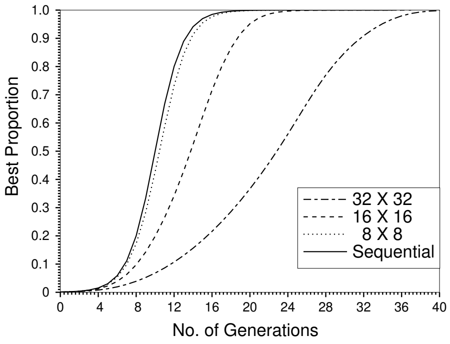

Fig. 7. Sequential and emergent growth curves for fitness proportional selection using different grid sizes and radius 0.89.

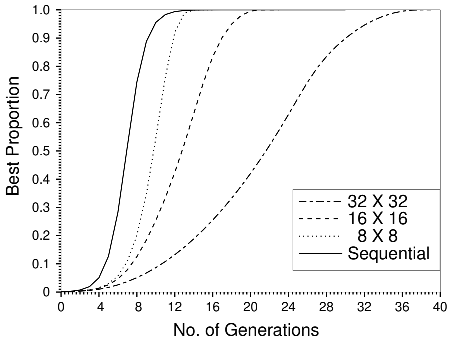

Fig. 8. Sequential and emergent growth curves for linear ranking selection using different grid sizes and radius 0.89.

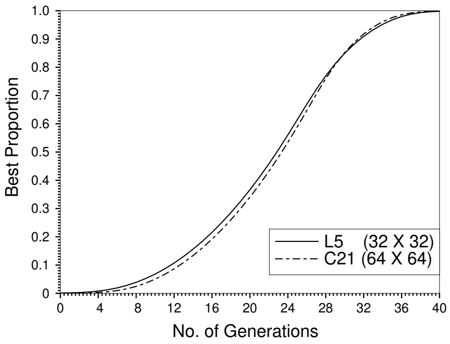

Fig. 9. Emergent growth curves for fitness proportional selection using constant radius ratios.

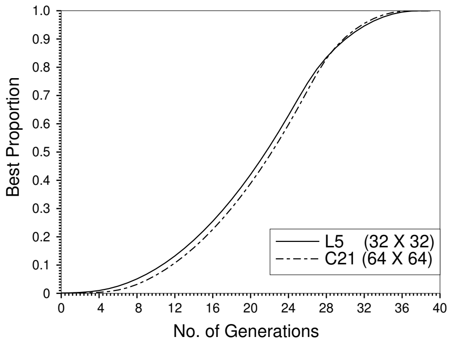

Fig. 10. Emergent growth curves for linear ranking selection using constant radius ratios.

local neighborhood models are well approximated by a logistic equation of the form:

$$N = \frac{N^*}{1 + (\frac{N^*}{N_0} - 1)e^{-at}}$$

where N^* is the asymptotic value of N for very large t and N_0 is the initial value of N [3]. Since in our case we know that N^* is 1.0 we can rewrite the equation as:

$$P_{b,t} = \frac{1}{1 + (\frac{1}{P_{b,0}} - 1)e^{-at}}$$

where a is the growth coefficient and $P_{b,t}$ is the proportion of the best in the population at time t. In this case the growth rate is controlled by a single parameter a, the coefficient in the exponent of the above equation.

In the case of sequential EAs the values of a is the natural logarithm of the fitness ratio for proportional selection and 1 for linear ranking [1, 6]. However, for spatially structured populations the growth coefficient a is dependent on the ratio of the neighborhood radius to the grid radius. It is fairly straight forward to plot the growth curves associated with various radius ratios in the interval $(0, 1)$ and estimate the associated value of a by using a least squares fit of the best logistic curve to the generated growth curves.

If we then plot the value of a as a function of increasing ratio r we obtain data of the sort illustrated in Figure 11. The perhaps not too surprising result is a very clear inverse exponential relationship between r and a with the case of $r = 1$ and the value of a corresponding to the sequential case.

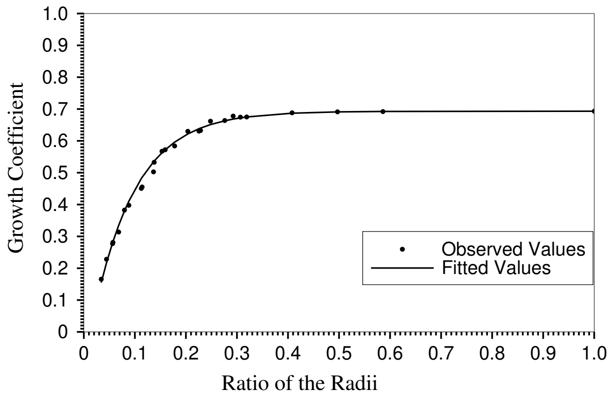

Fig. 11. Emergent growth coefficients for fitness proportional selection plotted against radii ratios.

As one might expect, selection intensities comparable to sequential EAs are achieved already with radii ratio (r) values of 0.5 (see Figure 11).

8 Discussion and Conclusions

Previous studies involving local neighborhood EAs selected the size and shape of local neighborhoods primarily on the basis of empirical tuning studies. The results presented in this paper provide a much clearer and more systematic way to make such decisions. It is shown that the critical parameter is the ratio of the radius of the neighborhood to the radius of the underlying grid. By varying this ratio the emergent selection intensity induced on the global population can be directly controlled. Moreover, a quantitative model of this relationship has been presented in which the coefficient of the growth rate of the best individual in the population is shown to be an inverse exponential function of this ratio.

A direct consequence of this analysis is the observation that typical neighborhood sizes and shapes used in the literature result in much lower selection pressures than their serial counterparts, but that comparable selection intensities can be obtained (if desired) by simply increasing this ratio to 0.5.

Effective EAs, whether serial or parallel, maintain a balance between exploration (via well chosen operators) and exploitation (well chosen selection pressure). Understanding how selection pressure can be directly controlled in local neighborhood EAs via a single parameter provides a simple means for adjusting that balance.

We believe that these observations will hold for other topologies and other local selection algorithms as well and are currently working on these issues.

References

1. Ankenbrandt, C. A. (1991). An extension to the theory of convergence and a proof of the time complexity of genetic algorithms. In *Foundations of Genetic Algorithm*, San Mateo, CA, pp. 53–68. Morgan Kaufmann.
2. Baluja, S. (1993). Structure and performance of fine-grained parallelism in genetic search. In *Proceedings of the Fifth International Conference on Genetic Algorithms*, Urbana–Champaign, IL, pp. 155–162. Morgan Kaufmann.
3. Banks, R. B. (1994). *Growth and Diffusion Phenomena*. Springer-Verlag.
4. Collins, R. J. and D. R. Jefferson (1991). Selection in massively parallel genetic algorithms. In *Proceedings of the Fourth International Conference on Genetic Algorithms*, San Diego, CA, pp. 249–256. Morgan Kaufmann.
5. De Jong, K. and J. Sarma (1995). On decentralizing selection algorithms. In *Proceedings of the Sixth International Conference on Genetic Algorithms*, Pittsburgh, PA, pp. 17–23. Morgan Kaufmann.
6. Goldberg, D. E. and K. Deb (1991). A comparative analysis of selection schemes used in genetic algorithms. In *Foundations of Genetic Algorithms*, San Mateo, CA, pp. 69–93. Morgan Kaufmann.
7. Gorges-Schleuter, M. (1989). ASPARGOS an asynchronous parallel genetic optimization strategy. In *Proceedings of the Third International Conference on Genetic Algorithms*, Fairfax, VA, pp. 422–427. Morgan Kaufman.
8. McInerney, J. (1992). *Biologically Influenced Algorithms and Parallelism in Non–linear Optimization*. Ph. D. thesis, University of California, San Diego.
9. Mühlenbein, H. (1991). Evolution in time and space – the parallel genetic algorithm. In *Foundations of Genetic Algorithm*, San Mateo, CA, pp. 316–337. Morgan Kaufmann.
10. Spiessens, P. and B. Manderick (1991). A massively parallel genetic algorithm: Implementation and first analysis. In *Proceedings of the Fourth International Conference on Genetic Algorithms*, San Diego, CA, pp. 279–287. Morgan Kaufmann.
11. Unwin, D. (1981). *Introductory Spatial Analysis*. Methuen and Co. Ltd.