The Spatially-Dispersed Genetic Algorithm

Grant Dick

Department of Information Science
University of Otago, PO Box 56, Dunedin, New Zealand
gdick@infoscience.otago.ac.nz

Abstract. Spatially structured population models improve the performance of genetic algorithms by assisting the selection scheme in maintaining diversity. A significant concern with these systems is that they need to be carefully configured in order to operate at their optimum. Failure to do so can often result in performance that is significantly under that of an equivalent non-spatial implementation. This paper introduces a GA that uses a population structure that requires no additional configuration. Early experimentation with this paradigm indicates that it is able to improve the searching abilities of the genetic algorithm on some problem domains.

1 The Spatially-Dispersed Genetic Algorithm

The spatially-dispersed genetic algorithm (sdGA) is an alternative method of incorporating population genetics models into genetic algorithms by using a two dimensional Euclidean space to hold the members of the population [1]. This space is infinite and continuous. The placing of individuals on a plane restricts them so that they can breed only with individuals that are "visible" to them.

In the sdGA, the first parent is selected from the entire population via whichever selection method the system is currently using. A second parent is then selected from the deme that is visible to the initial parent. This concept of visibility essentially creates subpopulations based on the location of individuals on the surface.

When offspring are created, they are placed in a location that maintains the relationship with their parents' coordinates. To introduce new offspring into the population, one parent is chosen randomly and the new individual is placed randomly at coordinates which fall within that parent's visibility radius.

The sdGA does not require significant tuning of parameters in order to achieve its maximum potential, since the spatial behaviour of the population rapidly becomes independent of the selected visibility and initial scale space [2]. This differs from other spatially-structured population strategies (for example, an Island GA), which often require extensive parameter tuning [3].

2 Experimental Results

The sdGA was tested on the Travelling Salesman Problem (TSP) and Holland's Royal Road function. Each problem was tested in a generational and steady-

E. Cantú-Paz et al. (Eds.): GECCO 2003, LNCS 2724, pp. 1572–1573, 2003.

state implementation. The results were compared with a genetic algorithm that used a non-spatially structured population.

The generational sdGA shows a significant improvement in performance over that of the generational standard GA for the TSP. Not only is the quality of the discovered solution better, the time taken to find an optimal solution is significantly reduced. In contrast, the steady-state sdGA appears to outperform the steady-state standard GA in the early iterations, but the advantage tapers out in the later stages of execution.

The Royal Road function appeared to behave in a completely opposite manner to the TSP. For the Royal Road problem, the steady-state sdGA performs better than the steady-state standard GA. The difference is discernible from as early as 2000 fitness evaluations. The generational sdGA appears to offer equal performance to that of the generational standard GA.

3 Conclusion and Future Work

The spatially-dispersed GA is a new technique for incorporating population genetics concepts into genetic algorithms. A significant advantage that the sdGA has over its predecessors is that it requires no user intervention or parameter setting to operate closer to its optimum level. Current results show that it is able to improve the performance of a GA on some problem domains.

3.1 Future Work

The sdGA is still in its infancy. More work needs to be done to properly understand the effect that its population structure has on GAs. Of particular interest would be an investigation into the population structure's effect on the schema theorem.

The ultimate goal of the sdGA is to use it in a distributed implementation. This would result in a parameterless distributed GA. Such a system should be of great interest to the GA community and the implications of a distributed implementation, in terms of computational overhead and scalability, is something that future work should investigate.

References

1. Dick, G., Whigham, P.: Spatially constrained selection in evolutionary computation. In: Australisia-Japan Joint Workshop on Intelligent and Evolving Systems. (2002)
2. Whigham, P., Dick, G.: A study of spatial distribution and evolution. In: The 14^{th} Annual Colloquium of the Spatial Information Research Centre, Wellington, New Zealand. (2002) 157–166
3. Cantú-Paz, E.: A survey of parallel genetic algorithms. Technical Report 97003, Illinois Genetic Algorithms Laboratory(IlliGAL) (1997)