Asynchronous Parallel Evolutionary Algorithms: Leveraging Heterogeneous Fitness Evaluation Times for Scalability and Elitist Parsimony Pressure

Matthew A. Martin Natural Computation Laboratory Department of Computer Science Missouri University of Science and Technology Rolla, Missouri, U.S.A. mam446@mst.edu Alex R. Bertels Natural Computation Laboratory Department of Computer Science Missouri University of Science and Technology Rolla, Missouri, U.S.A. arb9z4@mst.edu Daniel R. Tauritz Natural Computation Laboratory Department of Computer Science Missouri University of Science and Technology Rolla, Missouri, U.S.A. dtauritz@acm.org

ABSTRACT

Many important problem classes lead to large variations in fitness evaluation times, such as is often the case in Genetic Programming where the time complexity of executing one individual may differ greatly from that of another. Asynchronous Parallel Evolutionary Algorithms (APEAs) omit the generational synchronization step of traditional EAs which work in well-defined cycles. This paper provides an empirical analysis of the scalability improvements obtained by applying APEAs to such problem classes, aside from the speed-up caused merely by the removal of the synchronization step. APEAs exhibit bias towards individuals with shorter fitness evaluation times, because they propagate faster. This paper demonstrates how this bias can be leveraged in order to provide a unique type of "elitist" parsimony pressure which rewards more efficient solutions with equal solution quality.

Categories and Subject Descriptors

I.2.8 [Artificial Intelligence]: Problem Solving, Control Methods, and Search

Keywords

Parallelization, Asynchronous, Parsimony Pressure

1. INTRODUCTION

Evolutionary Algorithms (EAs) are inherently parallel due to their ability to simultaneously evaluate the fitness of individuals. Synchronous Parallel EAs (SPEAs) leverage this with the intent to gain significant speedups when executed on multiple processors. However, many important problem

GECCO '15 July 11-15, 2015, Madrid, Spain © 2015 Copyright held by the owner/author(s). ACM ISBN 978-1-4503-3488-4/15/07.

 ${\tt DOI: http://dx.doi.org/10.1145/2739482.2764718}$

classes lead to large variations in the computational time needed to evaluate the fitness of a trial solution. One example is when dealing with sampled fitness, where fitness may for instance be estimated by applying the trial solution to a set of instances sampled from a problem domain specific distribution; some instances may take significantly more computational time to evaluate than others. Another example is when the complexity of the genotypes in a population can significantly vary, such as is often the case in Genetic Programming (GP), where typically limits are placed on genotype size (tree depth in Koza style GP) and larger genotypes are penalized to create parsimony pressure to combat bloat [1, 3]. When attempting to achieve performance gains by parallelizing EAs through distributing the evaluation of individuals over slave nodes, while maintaining the synchronous nature of standard EAs by waiting for all individuals to be evaluated before commencing with survivor selection, such variations greatly reduce performance by causing idling until the last slave node has completed. In the extreme case of all n slave nodes but one, completing in a single unit of time, while one slave node takes tunits of time to complete, the percentage of wasted clock cycles in SPEAs is $\frac{(n-1)\cdot(t-1)}{n\cdot t}$, which tends towards 100% as n and t increase. Asynchronous Parallel EAs (APEAs) address this issue by performing survivor selection and all other evolutionary processes, without waiting for all individuals to be evaluated in the customary batch process. Since parallelizing EAs can be very time-consuming, many papers - including this one - rely on adding artificial evaluation time [2, 5, 6], which as a side-effect makes actual parallel execution unnecessary. Previous published work has mainly focused on the speed-up caused by the removal of the synchronization step in terms of less wasted clock cycles. In 2012, work by Yagoubi and Schoenauer [5] implies that the asynchronous model exhibits an implicit parsimony pressure that causes individuals with shorter evaluation time to be evolved, but this was not investigated.

2. METHODOLOGY

The experiment reported here demonstrates the presence of implicit parsimony pressure in APEAs, by explicitly encoding evaluation time into an individual's gene, which al-

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).



Figure 1: Box plot of average evaluation time of the final population of each run. The results are paired in increments of five slave nodes where blue indicates the results of the asynchronous runs and red indicates the results of the synchronous runs.

lows an arbitrary fitness function to simulate a wide range of problems whose actual evaluation time may be relatively long, such as GP. The Royal Road function [4] was chosen due to being challenging yet solvable, while fast to execute. An extra gene was inserted into all individuals' genotypes to represent the simulated evaluation time encoded as a floating point number. Uniform crossover is employed for this new gene, where offspring randomly receive one of the parents' evaluation times. The mutation operation is a Gaussian mutation with a mean of 0 and a standard deviation of 0.2. If the mutation causes the evaluation time to exceed the random initialization bounds, it is set to the bound it exceeded. This guarantees that encoding evaluation time cannot affect the performance of SPEAs. However, the hypothesis is that in APEAs it should create an "elitist" parsimony pressure causing the average evaluation time to decrease and thus increase the performance; the pressure should rise as the number of slave nodes increases. Where standard parsimony pressure generally can make EAs non-elitist, i.e., they can cause high-quality individuals to be discarded if they are slow running unless carefully tuned to avoid this, this implicit parsimony pressure cannot make an elitist APEA become non-elitist as selection is based strictly on fitness.

3. RESULTS & DISCUSSION

The results when compared to the experiments where the evaluation times were not encoded in the genes, show that there is a parsimony-like pressure acting on the population, but unlike standard parsimony pressure as it does not affect the fitness of the individual. This means the parsimony pressure is "elitist", as the ordering of the solutions does not change based on the evaluation time of the individual. To gain insight into the "elitist" parsimony pressure, the average evaluation times of the final population were analyzed versus the number of slave nodes. Figure 1 shows the average evaluation time of the individuals in the final population for each given run. These results are averaged over 100 runs for each number of slave nodes tested. As can be seen, there is a general trend down to roughly 1.5 seconds. From this graph it appears as though the parsimony pressure rises as the number of slave nodes is increased.

4. CONCLUSIONS & FUTURE WORK

This paper introduced an APEA simulation where artificial evaluation times are explicitly encoded in the genotype rather than randomly determined at evaluation time, thus allowing the APEA simulation to imitate a real APEA with far higher fidelity than in previous published work. It demonstrates through the use of this higher fidelity model, the potential for APEAs to exhibit a unique "elitist" parsimony pressure towards faster evaluating individuals, which rises as the number of slave nodes increases. Further research needs to be conducted in order to determine the degree to which this pressure manifests itself in unsimulated APEAs where the heterogeneous evaluation time is implicit. For example, in an unsimulated APEA, the actual evaluation time of an offspring may not be accurately reflected by the uniform crossover of the parents' explicitly encoded artificial evaluation times like employed in this research, and it is unclear to what extent this may impact the results.

The authors believe that the "elitist" parsimony pressure introduced by the asynchronous model will be similar in effectiveness to traditional parsimony pressure mechanisms, while having the benefit of being "elitist". Further research is needed to validate this. Also, it should be noted that it shares with traditional parsimony pressure [5] an increased difficulty in finding the optimal solution if the optimal solution has an extremely slow evaluation time.

5. REFERENCES

- S. Bleuler, M. Brack, L. Thiele, and E. Zitzler. Multiobjective Genetic Programming: Reducing Bloat Using SPEA2. In *Proceedings of the 2001 Congress on Evolutionary Computation*, volume 1, pages 536–543. IEEE, 2001.
- [2] A. W. Churchill, P. Husbands, and A. Philippides. Tool Sequence Optimization using Synchronous and Asynchronous Parallel Multi-Objective Evolutionary Algorithms with Heterogeneous Evaluations. In 2013 IEEE Congress on Evolutionary Computation (CEC), pages 2924–2931. IEEE, 2013.
- [3] J. Miller. What bloat? Cartesian Genetic Programming on Boolean problems. In 2001 Genetic and Evolutionary Computation Conference Late Breaking Papers, pages 295–302, 2001.
- [4] M. Mitchell, S. Forrest, and J. H. Holland. The Royal Road for Genetic Algorithms: Fitness Landscapes and GA Performance. In *Proceedings of the First European Conference on Artificial Life*, pages 245–254. Cambridge: The MIT Press, 1992.
- [5] M. Yagoubi and M. Schoenauer. Asynchronous Master/Slave MOEAs and Heterogeneous Evaluation Costs. In Proceedings of the Fourteenth International Conference on Genetic and Evolutionary Computation Conference, pages 1007–1014. ACM, 2012.
- [6] M. Yagoubi, L. Thobois, and M. Schoenauer. Asynchronous Evolutionary Multi-Objective Algorithms with Heterogeneous Evaluation Costs. In 2011 IEEE Congress on Evolutionary Computation (CEC), pages 21–28. IEEE, 2011.