

Guest editorial: special issue on parallel and distributed evolutionary algorithms, part two

Marco Tomassini · Leonardo Vanneschi

Published online: 7 April 2010
© Springer Science+Business Media, LLC 2010

Welcome to part 2 of the two-part special issue on parallel and distributed Evolutionary Algorithms (EAs). This special issue targets a wide set of researchers ranging from specialists in parallel computing for evolutionary computation to researchers interested in using the methodology for solving difficult practical problems. As already pointed out in the editorial of part 1 [1], despite the long history of parallel and distributed approaches in EAs, still a lot of motivations exist for a new presentation of the state of the art in the field, ranging from the fact that engineers and scientists wish to tackle problems of ever increasing complexity with a consequent increase in computational costs, to the growing availability of relatively cheap parallel/distributed hardware and the emergence of standard software environments (for instance MPI or the numerous existing frameworks for exploiting the computational power of computer GRIDs). Another remarkable emerging trend, well represented by the papers that appear in the present issue, consists in the use of graphical processing units (GPUs), which are inexpensive, can be easily harnessed, and provide excellent speedup in the case of evolutionary algorithms.

The fact that parallel and distributed EAs is still a very active research field is also testified by the unusually high number (twenty-five) of high-quality submissions that we received in response to our call for papers, ten of which were accepted for publication after extensive review. Five of these papers appeared in part 1 of this special issue [1], while the remaining five appear here. All of the papers were subject to the same review and editorial standards as any regular paper in this journal.

M. Tomassini (✉)
Information Systems Department, HEC, University of Lausanne, Lausanne, Switzerland
e-mail: marco.tomassini@unil.ch

L. Vanneschi
Department of Informatics, Systems and Communication (D.I.S.Co.),
University of Milano-Bicocca, Milan, Italy

In the first paper, Folino et al. present a distributed data mining algorithm to improve detection accuracy of malicious or unauthorized activity in a network context, including intrusion detection and cyber attacks. The approach is based on Genetic Programming (GP) and the distributed model is a hybrid that combines the cellular and the island model. Results on extended data sets compare favorably with other classification methods.

The second paper, by Wilson and Banzhaf, is another example of the present trend in GP towards the use of powerful and inexpensive GPUs for data-intensive parallel processing. They describe in great detail a parallel implementation of linear genetic programming on heterogeneous devices. Performance has been evaluated and all the design decisions needed to reach the best compromise are carefully discussed.

The third paper, by Hoverstad, apparently follows a different trend. Instead of parallelizing at the data level on local GPUs, for example, it addresses a coarser-grain model in which the target architecture is typically a cluster of standard machines. This is what has been traditionally done in the last decade, given the ubiquitous availability of this type of computing configuration. However, the contributed article describes a system which is publicly available, portable, and very easy to configure and use, which should make distributed EA processing easily accessible to a wide audience.

The fourth paper, by Hu et al., investigates the effect of a variable population size on the performances of an EA, an almost neglected aspect of evolutionary systems that could nevertheless be important to harness. Once again, the parallelization framework is that of GPUs. The authors conclude that an intrinsically variable population size may be helpful in improving performance with respect to the traditional fixed-size framework.

The final paper in this issue, by Laredo et al., deals with dynamical populations in P2P infrastructures in which small-world-like communication topologies play an important role. In particular the role played by these structures in preserving diversity has been studied using a family of trap functions. They found that the performance of an evolutionary algorithm using these topologies on this class of functions is superior to that of a standard setting, generally using a smaller population and thus less computational effort.

The guest editors thank all the authors who submitted their work for consideration. We are also much indebted to the many colleagues who kindly made their time available to assist in the reviewing process. Special thanks go to Lee Spector, the GPEM Editor-in-Chief, for his continuing assistance and encouragement during all the phases of the project. Finally, the issue could not have been possible without the dedication and professional work of the people at Springer, in particular Melissa Fearon, Emma Kalb, and Vijay Kumar.

Reference

1. *Genetic Programming and Evolvable Machines*, 10(4), December 2009