New Generation Computing, 29(2011)125-128
Ohmsha, Ltd. and Springer

Preface

Special Issue: Bio-Inspired Optimization Techniques for High Performance Computing

Gianluigi FOLINO and Carlo MASTROIANNI
ICAR-CNR, Via P. Bucci 41C, 87036 Rende(CS), ITALY
{folino,mastroianni}@icar.cnr.it

Received 14 January 2011

Keywords: Bio-inspired Algorithms, Optimization, Swarm Intelligence, Ant Colony Optimization, Parallel and Distributed Algorithms.

In the last few years, bio-inspired algorithms, mimicking the darwinian evolution or the behavior of ant colonies, flocks of birds, insect swarms,[1] etc.., have emerged as a viable solution to many parallel and distributed computational problems, and they have proved effective especially when distributed systems need adaptive and fault-tolerance properties.[3,4]

Their inherent parallelism and scalability make these kinds of algorithm very suitable for dynamically changing environments and systems, such as Grid Computing, Cloud Computing, and peer-to-peer systems, which are characterized by decentralized control, large scale and extreme dynamism. The availability of cheaper and better-performing distributed and parallel architectures has permitted on the one hand to foster the use of bio-inspired algorithms for these kinds of architecture, on the other hand to cope with complex problems, which were previously hardly solved owing to the large amount of needed time.

This special issue is dedicated to high performance bio-inspired algorithms for solving hard real world problems and to the design and implementation of distributed computing systems using this kind of algorithms. The issue collects the revised and extended versions of three papers selected among those presented at BADS 2010,[2] the 2nd edition of the Workshop on Bio-Inspired Algorithms for Distributed Systems that was hosted by ICAC 2010, the 7th IEEE International Conference on Autonomic Computing held in Washington, DC, USA, in June 2010.

The three papers were subjected to two more rounds of review with a minimum of three reviewers, which ensured the very high quality level of the research work published here. As for the application domain, two of these papers apply distributed techniques to very hard optimization problems, while the last paper uses an ant-based algorithm to give a novel contribution to the problem of mapping Virtual Machines to the servers of a Cloud platform. All the selected works are applied to topical real world domains as protein modeling applications, drug discovery and mapping of software components in distributed environments.

PSO (Particle Swarm Optimization) is an optimization technique[5)] inspired by the behavior of flocks of birds and shoals of fish. The parameter space is explored by a number of simple entities, the particles. Each particle evaluates a given fitness function at its current location, and then computes its new direction and position on the basis of its past best fitness value and that of its neighbors. PSO is becoming increasingly popular, and has applications in many domains, thanks to its effectiveness and extremely easy implementation. The work of Vanneschi et al. introduces four new parallel and distributed PSO methods and compare their performance on the well-known *Rastrigin* set of functions and on five complex real-life applications in the field of drug discovery. Specifically, the Particle Swarm Evolver (PSE) algorithm is a variant of genetic algorithms, in which each individual in the population is itself modeled as an evolving swarm. The Repulsive PSE is a variant in which a repulsive behavior is added, i.e., each particle is repulsed by the global best of all the other swarms in the GA population. Both multi-swarm PSO (MPSO) and Repulsive Multi-swarm PSO (RMPSO) make a PSO swarm evolve in parallel, replacing the k worst particles with the k best particles of the neighboring swarm. The main difference is that in RMPSO one half of the swarms have a repulsive behavior. Experiments show that RMPSO outperforms the other considered PSO methods on all the analyzed problems, probably due to the fact that it is able to maintain a higher diversity degree in the whole system.

Zhu et al. exploit the potentialities of Graphic Processing Units (GPUs) to cope with hard multi-objective optimization problems. The authors introduce the Differential Evolutionary Markov Chain Monte Carlo (DEMCMC) method, which combines Differential Evolution (DE) and Markov Chain Monte Carlo (MCMC) advantages to make a population of Markov chains evolve toward a diversified set of solutions at the Pareto optimal front. Then, they describe an extension of this algorithm, named DEMCMC-GPU, which adopts the Single Instruction Multiple Thread (SIMT) programming model for an efficient implementation on GPUs. The new algorithm is based on a master-slave paradigm, where the slave threads in the GPU carry out the computations for the Markov chains evolution, since a large population size is necessary to achieve a sufficient coverage of the Pareto optimal front. Performance evaluation is conducted on an NVIDIA GeForce GTX 480 GPU with 480 SIMT processors using standard benchmark functions (the DTLZ test suite) and an application for modeling protein loop structures. The adoption of an efficient parallel pseudo-random number generator along with a cache mechanism and the use of a large number of threads permitted to decrease the execution time by a factor of 100 in comparison with the CPU-only implementation.

The paper by Csorba et al. copes with a very important problem in distributed systems, i.e., the efficient mapping of software components to the nodes of a network. This problem is known to be NP-hard, so efficient heuristics are needed to find near optimal solutions in short time and in harsh network conditions. The paper presents a bio-inspired optimization framework built around the Cross Entropy Ant System (CEAS), a technique derived from Ant Colony Optimization. The framework uses ant-like agents that move around in the network, identify potential locations where components might be placed, and leave pheromone trails as a means to facilitate indirect communication among agents. While the approach is general, it must be adapted to the different application domains through the definition of cost functions that take into account the specific QoS requirements and the constraints of the particular problem. A significant merit of this study is that the approach is evaluated in three real application scenarios, that is: the clustering of collaborating software components, the deployment of component replicas in an environment with strict dependability requirements, and the placement of virtual machines in a private/public cloud computing scenario, where different clouds are associated with specific financial costs.

The papers included in this special issue confirm that bio-inspired algorithms can not only adapt to work in novel high performance (Grid, Cloud, GPU) environments, but they can also provide characteristics as decentralization, self-organization, flexibility and energy saving that are essential to efficiently cope with the ever increasing complexity of these systems.

We hope that you will find this special issue interesting, useful, and informative and it will give you a valuable inspiration for your research. We are sure you will enjoy reading these articles as much as we did. In conclusion, we would like to thank the Program Committee members and the reviewers of BADS'10 for their valuable input, insight, and expert comments, and the Editor-in-Chief, Dr. Taisuke Sato, for his valuable advice and strong support during the preparation of this special issue.

References

1) Dorigo, M., Bonabeau, E. and Theraulaz, G., "Ant algorithms and stigmergy," *Future Generation Computer Systems, 16, 9*, pp. 851–871, 2000.

2) Folino, G., Mastroianni, C., Fragopoulou, P. and Suzuki, J. eds., *2nd Workshop on Bio-Inspired and Self-* Algorithms for Distributed Systems*, ACM, June 2010.

3) Folino, G. and Spezzano, G., "An autonomic tool for building self organizing grid-enabled applications," *Future Generation Comp. Syst., 23, 5*, pp. 671–679, 2007.

4) Forestiero, A., Mastroianni, C. and Spezzano, G., "So-grid: A self organizing grid featuring bio-inspired algorithms," *TAAS, 3, 2*, 2008.

5) Kennedy, J. and Eberhart, R., "Particle swarm optimization," in *Neural Networks, 1995. Proc., IEEE International Conference on, 4*, pp. 1942–1948, August 2002.

Gianluigi Folino, Ph.D.: He received the M.Sc. degree in engineering in computer science from the University of Calabria, Cosenza, Italy, in 1997 and the Ph.D. degree in computer science from Radboud University, Nijmegen, Nijmegen, The Netherlands, in 2010. He joined ICAR-CNR, Institute of High Performance Computing and Networking, Italian National Research Council, in 1999, supported by the National Institute of Nuclear Physics (INFN) fellowship. Since 2001, he has been a Researcher at ICAR-CNR, Cosenza, Italy. Since 2001, he has also been a Lecturer at the University of Calabria, Arcavacata di Rende, Italy, and at the University Magna Graecia of Catanzaro, Catanzaro, Italy, for a number of courses (operating system laboratory, operating systems and networks, foundations of computer science, and applied computer science). He was a Visiting Researcher at the University of Nottingham, U.K., in 2007, and at Radbound University, Nijmegen, the Netherlands, in 2008 and 2009. His current research interests include applications of distributed computing in the area of cellular automata, genetic programming, swarm intelligence, data mining, and bioinformatics.

Carlo Mastroianni, Ph.D.: He received the Laurea degree and the Ph.D. degree in Computer Engineering from the University of Calabria, Italy, in 1995 and 1999, respectively. He is a Researcher at the Institute of High Performance Computing and Networking of the Italian National Research Council, ICAR-CNR, in Cosenza, Italy, since 2001. Previously, he worked at the Computer Department of the Prime Minister's Office, in Rome. He co-authored over 80 papers published in international conference proceedings and journals, among which IEEE/ACM Transactions on Networking, ACM Transactions on Autonomous and Adaptive Systems, IEEE Transactions on Evolutionary Computations, Parallel Computing. He edited special issues for the journals Future Generation Computer Systems, Journal of Network and Computer Applications, Multiagent and Grid Systems. His areas of interest are Grid and Cloud computing, P2P networks, bio-inspired algorithms and protocols, distributed data mining.