

Special Section: Bio-Inspired Algorithms for Distributed Systems

Gianluigi Folino, Carlo Mastroianni

*ICAR-CNR, Via P. Bucci 41C, 87036 Rende(CS), Italy
{folino,mastroianni}@icar.cnr.it*

Abstract

This special section is dedicated to the use and evaluation of bio-inspired algorithms for the design and implementation of distributed computing systems. This issue collects the revised and extended versions of the five best papers presented at BADS 2009, the Workshop on Bio-Inspired Algorithms for Distributed Systems that was hosted by ICAC 2009, the 6th IEEE International Conference on Autonomic Computing held in Barcelona, Spain, in June 2009. The papers of this special section present bio-inspired solutions for important problems such as overlay construction and resource discovery in P2P networks, job mapping in a heterogeneous environment, and data dissemination and aggregation in wireless sensor networks, for which special attention is given to the important issue of energy saving. The five papers confirm that the bio-inspired paradigm naturally provides such characteristics as decentralization, self-organization, flexibility, and energy saving, which are essential to efficiently cope with the ever increasing complexity of distributed computing systems.

Computer systems are characterized by an ever growing complexity and a pronounced distributed nature. While the use of centralized or hierarchical architectures and algorithms has been dominant so far, they are now becoming impractical because of their poor scalability and fault-tolerance characteristics. Decentralized architectures, for example P2P, Grid and Cloud systems, are increasingly popular, but they need new types of algorithms to be managed efficiently.

The study of natural ecosystems helped to understand how a complex system can easily adapt to a changing environment, not through the deployment of a central strategy, but exploiting the evolving and adapting behavior of a multitude of actors, and the cooperation among them. Examples come from a wide range of biological systems, including ant colonies, bird flocking, honey bees, bacteria, and many more (Dorigo et al., 2000). The solution of a problem can emerge from the activity of “intelligent” agents that perform complex functionalities, or from the interaction of a large number of very simple agents, in the so called “swarm intelligence” systems. Such techniques are sometimes “evolutionary”, as they can exploit genetic rules for the selection and the recombination of candidate solutions.

In recent years, following the examples offered by nature, bio-inspired algorithms are emerging as a viable solutions to many parallel and distributed computational

problems, and they are proving capable of reinforcing the adaptive and fault-tolerance characteristics of distributed systems (Ko et al., 2008; Folino and Spezzano, 2007). These solutions are generally inspired by the behavior of flocks of birds, insect swarms, and, more often, of ant colonies (Forestiero et al., 2008; Pavani and Waldman, 2009).

The workshop on Bio-inspired Algorithms for Distributed Systems, BADS 2009 (Folino et al., 2009), was organized during the ICAC 2009 conference (Barcelona, June 2009) to offer researchers operating in this field the opportunity to promote their solutions and share their ideas in an interactive context. The success of the workshop confirmed that distributed computing systems are a very natural application field for bio-inspired algorithms and protocols, and the growing interest of many and renowned researchers seems to anticipate that several complex problems will find efficient solutions in the next years. The authors of the best five papers of the BADS workshop, chosen on the basis of reviewers’ scores, were invited to revise and extend their papers, and submit them to this special section. The papers underwent two more review rounds, which ensured the very high quality level of the research work published here.

The papers of this special issue present and assess a number of interesting solutions, which are diverse for

the problems that they tackle, the application domain for which they are devised, and the type of inspiration taken from nature. The papers present bio-inspired solutions for important problems such as overlay construction and resource discovery in P2P networks, task/node mapping and scheduling in a heterogeneous environment, data aggregation in distributed systems, data dissemination in wireless sensor networks. The presented algorithms are tailored for specific domains, including P2P and Grid systems, and also rapidly evolving infrastructures such as Clouds and sensor networks. Interestingly, though, they can be easily adapted to domains different from those for which they have been devised, which can be a source of inspiration for the readers of this special section.

The presented solutions are also diverse with respect to the kind of inspiration taken from nature. The papers by Brocco et al., and by Biccocchi et al., are directly inspired by phenomena observed in ant colonies: the self-organizing constructions of topologies, and the mechanism of pheromone evaporation in ant foraging. Moreover, both these papers, and also the paper by Garbinato et al., take inspiration from epidemic mechanisms that help information to be rapidly disseminated over a distributed system (Guéret et al., 2007). The latter paper also exploits a power-law tuning of transmission power to reduce energy consumption in a sensor network, which is an interesting response to the emerging necessity to devise energy-aware solutions.

On the other hand, the papers by De Falco et al., and by González et al., take inspiration from the evolutionary behavior of natural systems, and focus on two variants of evolutionary algorithms, namely, Differential Evolution algorithms, and Parallel Genetic Programming (GP). While the former paper uses this type of approach to solve a specific problem - i.e., task/resource mapping - the latter offers an interesting analysis of the inherent fault-tolerance characteristics of Parallel GP, which confirms that algorithms of this kind can be used in a variety of dynamic domains.

The rest of this editorial piece offers a more detailed outlook on the five papers of this special section.

An important problem in distributed systems is the ability to share and discover resources within a community, for example data items in P2P systems or hardware/software resources in Grid infrastructures. Scalability and minimal network overhead are among the main issues to be addressed in these scenarios. The paper by Brocco et al. presents and evaluates a bio-inspired algorithm that aims to reduce the overall network traffic in a Grid. Ant-inspired mobile agents move across the Grid through P2P interconnections, and add

or remove logical links in order to optimize the network topology and maintain a bounded diameter overlay. The paper also introduces an effective resource discovery algorithm that executes over the obtained topology and uses cache information exchanged among the nodes via a gossiping mechanism. Experiment results reported in the paper prove that the presented approach allows discovery operations to be performed effectively, in a limited time and with low network overhead.

Desktop Grids (Kondo et al., 2007) provide large distributed computing resources when idle, at a very low cost. However, when using these large-scale heterogeneous platforms, the risk of failures is generally high. The paper by González et al. analyzes the behavior of Parallel Genetic Programming (PGP) applications executing in distributed platforms with high failure rates, with the goal of characterizing the inherent fault tolerance capabilities of the PGP paradigm. Two well-known GP problems, “even 5-parity” and “11-multiplexer”, are examined via simulation by using trace data collected from real-world Desktop Grid platforms. Experiments show that the PGP paradigm exhibits satisfying fault tolerance properties even in very dynamic environments. Indeed, PGP inherently provides *graceful degradation* without the need for additional fault tolerance techniques.

An efficient use of resources in a Grid environment requires an optimal task/node mapping. This is known as an NP-complete problem, therefore heuristic techniques must be adopted to find near-optimal solutions. Further degrees of complexity derive from the existence of user-dependent requirements - specified through QoS constraints - and from the time changing load of the involved resources. De Falco et al. propose the use of a multi-objective DE algorithm that is able to find near-optimal solutions to the mapping problem in a heterogeneous Grid environment, addressing both the user requirements and those of the Grid manager. The authors extend the classical Differential Evolution paradigm to cope with this problem, also by introducing a novel mutation operator. The algorithm is defined, implemented and evaluated in several allocation scenarios that differ by the types of application requests and the Grid operating conditions.

The paper by Garbinato et al. presents and evaluates a novel approach to decrease the power consumption of epidemic information dissemination in sensor networks. The adopted strategy consists in modulating the transmission range of sensors by correlating the energy used for transmission to a power-law statistical distribution. This strategy replicates the behavior of scale-free networks, since many nodes can reach few neigh-

bors, while few nodes can reach many neighbors. To evaluate their approach, the authors injected this power-law modulation into four existing epidemic algorithms and compared the related performances with the original versions, in which the transmission range is fixed or distributed uniformly. This evaluation showed that the power-law approach improves the efficiency of the original algorithms in terms of power consumption, with no negative impact on their effectiveness, measured in terms of the number of nodes that can be reached after the dissemination.

The paper by Biccocchi et al. tackles the important issue of deriving aggregation information about the properties of a large-size dynamic network - for example, the average load or the current number of nodes - and making it available at every node that needs it. The approach proposed by the authors is inspired by the physical/biological of *diffusion*, and is realized via *gossip*-based communication between the nodes of the network. To cope with dynamic scenarios, in which the aggregated values change continuously, the authors present a novel technique inspired by the mechanism of pheromone evaporation in ant foraging. In analogy with what happens to pheromone trails used by ants to discover food sources, the aggregated values computed with the diffusion strategy are made evaporate slowly so as to enforce updated information and globally obtain more accurate values. The method avoids the need for periodic restarts used by “epoch-based” gossiping techniques.

The papers included in this special session confirm that bio-inspired algorithms naturally provide such characteristics as decentralization, self-organization, flexibility and energy saving that are essential to efficiently cope with the ever increasing complexity of current and future distributed systems. We hope that you will find this special issue interesting and thought-provoking and that it will offer you an opportunity to deepen your knowledge on this field, and a valuable inspiration for your research.

In conclusion, we would like to thank the editors of FGCS and the people that helped us in the publication process, in particular, Ella Chen and Fred Kop. We are also grateful to the Program Committee members of BADS’09 and to the reviewers that helped us to select and polish the five papers.

References

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

Folino, G., Krasnogor, N., Mastroianni, C., Zambonelli, F. (Eds.), June 2009. BADS’09: Proceedings of the Workshop on Bio-Inspired Algorithms for Distributed Systems. ACM Press, Barcelona, Spain.
URL <http://bads.icar.cnr.it>

Folino, G., Spezzano, G., 2007. An autonomic tool for building self-organizing grid-enabled applications. *Future Generation Computer Systems* 23 (5), 671–679.

Forestiero, A., Mastroianni, C., Spezzano, G., May 2008. So-Grid: A self-organizing grid featuring bio-inspired algorithms. *ACM Transactions on Autonomous and Adaptive Systems* 3 (2).

Guéret, C., Monmarché, N., Slimane, M., 2007. A biology-inspired model for the automatic dissemination of information in P2P networks. *Multiagent and Grid Systems* 3 (1), 87–104.

Ko, S. Y., Gupta, I., Jo, Y., 2008. A new class of nature-inspired algorithms for self-adaptive peer-to-peer computing. *ACM Transactions on Autonomous and Adaptive Systems* 3 (3), 1–34.

Kondo, D., Fedak, G., Cappello, F., Chien, A. A., Casanova, H., 2007. Characterizing resource availability in enterprise desktop grids. *Future Generation Comp. Syst.* 23 (7), 888–903.

Pavani, G. S., Waldman, H., March 2009. Co-scheduling in lambda grid systems by means of ant colony optimization. *Future Generation Computer Systems* 5 (3), 257–265.