

Building Blocks from Biology for the Design of Algorithms for the Management of Modern Dynamic Networks

by Gianni A. Di Caro, Frederick Ducatelle, Luca Maria Gambardella, and Andrea Rizzoli

Modern computer and communication networks are becoming increasingly large, heterogeneous and dynamic. Traditional network algorithms fail to deal efficiently with this increased complexity. The EU-funded project BISON addresses this problem by drawing inspiration from biology to provide the building blocks for a new family of distributed, self-organizing, adaptive and scalable network algorithms.

Biology-Inspired techniques for Self-Organization in dynamic Networks (BISON) is a three-year Shared-Cost RTD Project (IST-2001-38923) funded by the Future and Emerging Technologies activity of the Information Society Technologies Program of the European Commission. It runs from January 2003 until December 2005. The BISON consortium draws on multidisciplinary expertise from the University of Bologna (Italy), which is the coordinator, Dalle Molle Institute for Artificial Intelligence (IDSIA) (Switzerland), the Technical University of Dresden (Germany) and Telenor AS (Norway), which is the industrial partner.

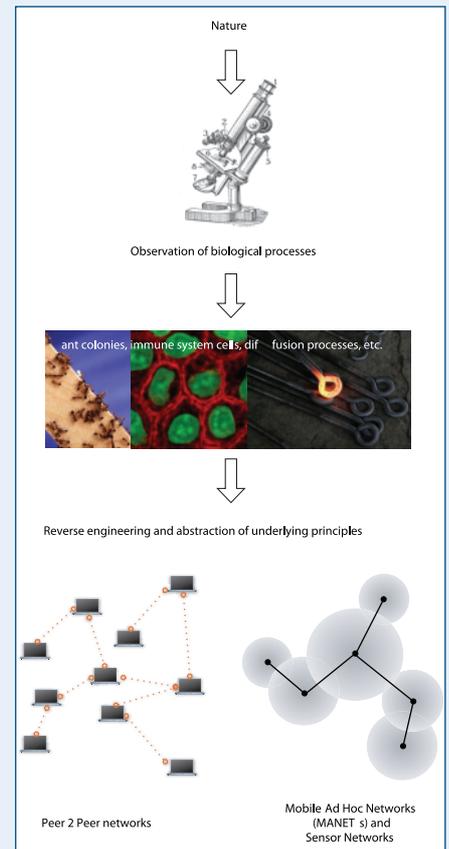
The objective of BISON is to develop new network algorithms that are adaptive to changes, robust to failures and perturbations, work in a self-organized and decentralized way, and are able to function efficiently in heterogeneous large-scale systems. Central to BISON's approach is the idea of drawing inspiration from biological systems to develop this new family of algorithms. The rationale behind this choice comes from the following observations. Biological systems usually have the ability to effectively adapt to constantly changing environments. Moreover, they are usually robust to internal perturbations or loss of units, and are able to survive and evolve in a wide range of environments.

Biological systems have obtained these properties through evolution: they are usually composed of a large number of dynamic, autonomous and distributed units, which display effective adaptive

behaviour at the system level as a result of local interactions and self-organization. Examples of systems showing this highly complex behaviour are ant colonies, and the cells of the human immune system. Because of these appealing properties, natural systems have served as a source of inspiration for a number of successful algorithms and frameworks, mainly for optimization tasks (eg evolutionary computation and ant colony optimization). The application of the biological approach to network problems has attracted relatively little attention. However, the analogy between networks and natural systems is evident: the nodes of a network environment can be seen as the units of a biological system, and the actions and input/output of users form the external environment.

The BISON research intends to exploit this analogy. We investigate biological processes in a systematic way, and abstract and reverse-engineer the basic mechanisms at work. This allows us to identify building blocks for the design of distributed and self-organizing network algorithms displaying adaptivity, robustness and scalability. BISON has focused on a number of specific network environments and core functions. Peer-to-peer (P2P) and multi-hop wireless ad hoc (MWAH) networks are the dynamic environments being investigated. The functions are routing, topology management, content search, monitoring, data aggregation and load balancing.

At IDSIA, we have focused on two networks. These are mobile ad hoc networks



From Biology to Dynamic Networks.

(MANETs), which are WMAH networks in which the nodes are mobile and can join or leave the network at any time, and sensor networks, in which nodes are usually not mobile and have specific sensing capabilities. In the case of MANETs, we addressed the issue of the optimization of the routing function to allow the network to cope effectively with the problems raised by mobility and interference. For this purpose we developed AntHocNet, a traffic- and topology-adaptive algorithm with both reactive and proactive components. AntHocNet is designed after ant colonies and their ability to find the shortest paths in distributed and dynamic environments by using a combination of repeated and concurrent path sampling, pheromone laying/following (stigmergic communication), and stochastic decisions.

For sensor networks we focused on the problem of the distributed assignment of node transmission ranges. Here the aim was to establish topologies that can provide full connectivity while minimizing energy consumption, thereby maximizing network lifetime. For sensor networks we developed both an exact cen-

tralized approach and an effective distributed heuristic based on a reaction-diffusion model, which is characteristic of a number of biological processes (eg morphogenetic development). Still in the context of ad hoc networks, but wired ones, we also developed a distributed active monitoring system based on the same ant colony behavior at the heart of AntHocNet.

In P2P networks, all nodes or users can realize bidirectional and symmetric communications. An overlay connecting the peers is established above the IP layer, which provides the underlying communication functionalities. Important issues in P2P networks include strategies for performing general network-wide calculations and searches, as well as building and maintaining overlay topologies. These can facilitate critical operations like content search/publishing and joining/leaving. In BISON we successfully addressed most of these issues.

Cell replication mechanisms provided inspiration for a protocol that builds and maintains random overlay topologies over time, while ideas from cell adhesion

are behind T-Man, a protocol provably able to build a variety of complex structured topologies. The pattern-matching and cell proliferation characteristics of antibodies of the human immune system have provided basic inspiration for an algorithm that effectively performs content searching in unstructured overlays. Diffusion processes and chemotaxis, a process describing cell movement in response to concentration gradients of chemicals, are at the foundations of a system for load balancing in distributed storage networks. A new proactive protocol for the calculation of aggregate functions (eg average load) has been derived from the patterns common to the epidemic spreading of contagious diseases.

In conclusion, the observation of biological processes has provided numerous ideas for the design of a number of fully distributed, adaptive, robust and scalable network algorithms. The experimental results from extensive simulations provide a strong validation of the approaches followed. Performance is usually good, and in some cases the proposed algorithms clearly outperform state-of-the-art algorithms. (This is the case of the algo-

rithms for routing and topology control in wireless networks.) In other cases, such as the protocols for aggregate calculations, excellent empirical performance is also accompanied by pleasing theoretical properties of convergence.

Links:

BISON web site:

<http://www.cs.unibo.it/bison>

IDSIA web site: <http://www.idsia.ch>

Ant Colony Optimization:

<http://iridia.ulb.ac.be/~mdorigo/ACO/ACO.html>

Artificial Immune Systems:

<http://www.dca.fee.unicamp.br/~lnunes/immune.html>

IETF MANET page: <http://www.ietf.org/html.charters/manet-charter.html>

Sensor networks: <http://www.research.rutgers.edu/~mini/sensornetworks.html>

P2P networks: <http://p2p.internet2.edu/>

Please contact:

Gianni A. Di Caro, Frederick Ducatelle, Luca Maria Gambardella, Andrea Rizzoli
Istituto "Dalle Molle" di Studi Sull'Intelligenza Artificiale (IDSIA), Manno-Lugano, Switzerland
Tel: +41 58 6666660

E-mail:

{gianni,frederick,andrea,luca}@idsia.ch

Chemical Programming of Self-Organizing Systems

by Jean-Pierre Banâtre, Pascal Fradet and Yann Radenac

Chemical programming relies on the 'chemical metaphor': data are seen as molecules and computations as chemical reactions. This programming paradigm exhibits self-organizing properties and allows the description of autonomic systems in an elegant way.

The Chemical Programming Model

This formalism was proposed to capture the intuition of computation as the global evolution of a freely interacting collection of atomic values. It can be introduced through the chemical reaction metaphor. The unique data structure is the multiset (a set possibly containing identical elements), which can be seen as a chemical solution. A simple program consists of a reaction condition and an action. Execution proceeds by replacing elements that satisfy the reaction condition with the elements specified by the action. The result of such a program is obtained when a stable state is

reached; that is to say, when no more reactions can take place. For example, the computation of the maximum element of a non-empty set can be described as:

replace x, y by x if $x \geq y$

Any couple of elements x and y of the multiset is replaced by x if the condition is fulfilled. This process continues until only the maximum element remains. Note that in this definition, nothing is said about the order of evaluation of the comparisons. If several disjoint pairs of elements satisfy the condition, reactions can be performed in parallel.

Chemical Programming and Self-Organizing Systems

Autonomic computing provides a vision in which systems manage themselves according to some predefined goals. The essence of autonomic computing is self-organization. Like biological systems, autonomic systems maintain and adjust their operation in the face of changing components, workloads, demands and external conditions, such as hardware or software failures, either innocent or malicious. The autonomic system might continually monitor its own use and check for component upgrades. We believe that