

# BISON: Biology-Inspired techniques for Self-Organization in dynamic Networks

Gianni A. Di Caro, Frederick Ducatelle, Luca M. Gambardella

*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 initiative of the Information Society Technologies Programme of the European Commission. It runs from January 2003 until December 2005.<sup>1</sup> The objective of BISON is the design and implementation of algorithms for management and control in modern network environments. These environments are typically large, complex, highly dynamic, and heterogeneous. To cope with these challenges, network algorithms should have fully distributed and self-organizing architectures, and be adaptive, robust, and scalable. The central claim of BISON is that the observation and analysis of complex adaptive systems found in Nature can be a major source of inspiration to design algorithms with such characteristics. In what follows, we first describe the objectives of the project. Then we present the specific classes of networks that have been investigated, and discuss some of the results obtained so far in conjunction with the biological mechanisms that have provided useful design guidelines. More details can be found at <http://www.cs.unibo.it/bison>.

## 1 Project objectives

Network systems are becoming increasingly complex. This is due to the size and heterogeneity of such systems, as well as to the complex interactions among their elements. Moreover, modern networks are often highly dynamic, with new services and contents continually made available, and nodes/users constantly entering/leaving the network and/or moving within it. Classical techniques for network management and control were not conceived to efficiently face such challenges. New algorithms are needed, which are adaptive to changes and robust to failures and perturbations, work in a self-organized and decentralized way, and are able to cope with heterogeneous large-scale systems. BISON proposes to draw inspiration from biological systems and the basic mechanisms regulating their functioning to develop this new family of algorithms. The rationale behind this choice comes from the following observations. Biological systems usually show the ability to effectively adapt to constantly changing environments. A majority of these systems can be seen as composed of a large number of dynamic, autonomous, and distributed units which through local interactions generate effective adaptive behaviors at the system-level as a result of self-organization. Moreover, biological systems are usually

<sup>1</sup> The BISON consortium is composed of: University of Bologna (Italy) which is the coordinator, IDSIA (Switzerland), Technical University of Dresden (Germany), Telenor AS (Norway).

robust to internal perturbations or loss of units, and are able to survive and evolve in wide ranges of environments. Looking at the nodes of a network environment as the units of a biological system and the users as the external environment, the analogy between the two classes of systems is clear.

Nature has already served as a source of inspiration for a number of successful algorithms and frameworks (e.g., evolutionary computation [6], ant colony optimization [4]). The aim of BISON is to investigate biological processes in a more systematic way and abstract and reverse-engineer the basic mechanisms at work in order to identify building blocks to design distributed and self-organizing network algorithms featuring adaptivity, robustness, and scalability.

The BISON effort is in the first place focused on a number of specific dynamic network environments and core functions. The investigated networks are wireless multi-hop ad hoc networks (WMAHNs) and peer-to-peer (P2P) networks. The functions are routing, topology management, content search, monitoring, data aggregation, and load balancing. The experience obtained using inspiration from biology to develop competitive algorithms for these tasks will eventually serve a more ambitious goal: to extract essential building blocks from biological systems and provide a framework for using them to construct solutions for dynamic networks.

## 2 Network environments

### 2.1 Wireless multi-hop ad hoc networks

Wireless multi-hop ad hoc networks are networks which are set up on-the-fly, using a set of wireless devices. These networks have no fixed infrastructure. All nodes are functionally equal and there is no central control or overview. The topology of the network is defined by the physical location of the nodes and the range of their radio transmitter. It can change constantly due to nodes joining or leaving the network, because nodes change their transmission range (adapting their transmission power), or possibly because the nodes are mobile. There are no designated routers in the network: all nodes serve as routers for each other, and data packets are forwarded from node to node in a multi-hop fashion. WMAHNs typically have low bandwidth due to the use of wireless communication and the need for complex medium access control mechanisms. WMAHNs in which nodes are mobile are called *Mobile Ad Hoc Networks* (MANETs) [17]. *Sensor Networks* [1] are an example of WMAHNs where nodes are usually not mobile. In BISON we have considered both these classes of networks, focusing on the optimization of the *routing function* for MANETs and on the *regulation of the radio range* for topology control and energy minimiza-

tion in sensor networks. We have also considered *monitoring* and *quality-of-service routing* in wired multi-hop ad hoc networks, which can be seen as a special case of WMAHNS. The latter work is not presented here; see [9] for details.

## 2.2 Peer-to-peer networks

Peer-to-peer networks are characterized by the fact that all nodes in the network act as “peers” and can realize bidirectional and symmetric communications. In a P2P network users/nodes can join and leave the network at any time and can directly exchange information through a fully distributed and self-organizing process. In practical terms, in P2P networks an *overlay network* connecting the communicating peers is established above the IP layer through end-to-end TCP/IP connections. That is, a virtual network is established at session/user level rather than at the physical/node level. The IP layer provides the basic communication functionalities, enabling IP-capable nodes to reach any other node in the IP network. On the other hand, the IP layer does not provide any information about where to find stored content or other participants to the network. This information must be somehow available at the level of the virtual overlay network connecting the peers. Different P2P networks rely on overlays with different topological characteristics and make use of different strategies to join, search, and publish/advertise. In BISON, we have focused our research on the virtual network level. A number of algorithms have been developed for calculating network-wide measures, content search and topology management. Only the first two of these topics will be discussed here. A detailed description of algorithms (inspired by *cell replication*) that can maintain a random topology can be found in [11], while [10] describes the protocol *T-Man*, which is able to build a variety of complex structured topologies using basic ideas from *cell adhesion*. We are also working on load balancing using diffusion equations mimicking a biological process called *chemotaxis*, in which cells move in response to concentration gradients of chemicals, but this work is still under validation.

## 3 BISON algorithms

In BISON a number of management and control functions for WMAHNS and P2P networks were investigated. For each of these functions algorithms were developed, taking inspiration from biological mechanisms. The algorithms were compared to state-of-the-art algorithms according to common performance measures, as well as measures which take into account the specific challenges of the considered dynamic networks, such as adaptivity, scalability and robustness.

### 3.1 Algorithms for multi-hop ad hoc nets

#### 3.1.1 Routing in mobile ad hoc networks

Routing is the task of directing data flows from sources to destinations maximizing network performance. It is at the very core of network functioning. In MANETs routing is

particularly challenging due to the dynamic aspects of the network and its bandwidth limitations. To face these challenges, we have developed a novel algorithm, called *AntHocNet* [3, 5]. The inspiration for the design of AntHocNet comes from the *foraging behavior of ant colonies*. It has been observed that ants in a colony can converge on moving over the shortest path connecting their nest to a food source [2]. The main catalyst of this colony-level shortest path behavior is the use of a chemical substance called *pheromone*. Ants moving back and forth between nest and food deposit pheromone, and preferentially move towards areas of higher pheromone intensity. Shorter paths can be completed quicker and more frequently by the ants, and are therefore marked with more pheromone. These paths attract more ants, which in turn increase the pheromone level. The overall effect is a *distributed reinforcement learning* process which eventually allows the majority of the ants to converge onto the shortest path. This behavior has attracted attention as a framework for optimization, and has been reverse-engineered in the *Ant Colony Optimization* metaheuristic [4].

AntHocNet is based on the same principles, thoroughly adapted to the challenges of MANETs. Nodes obtain routing information using ant-like agents which repeatedly *sample* and *reinforce* good paths. Routing information is maintained in arrays of pheromone variables, called *pheromone tables*, which play the role of routing tables: they are used to guide the ant agents searching for a path and in turn are updated to reflect the goodness (in terms of number of hops and traffic congestion) of the followed paths. AntHocNet contains both reactive and proactive components: it is *reactive* in the sense that nodes only gather routing information for destinations which they are currently communicating with, while it is *proactive* because nodes try to maintain and improve routing information as long as communication is going on. Such a hybrid architecture tries to combine the best of two worlds: while reactive algorithms are usually more efficient in terms of control overhead, proactive behavior allows better adaptivity. The proactive actions are realized through a combination of a *diffusion process* and the ant path sampling: the routing information maintained in the pheromone tables is spread at low frequency among the nodes of the MANET and is used both to incrementally update the pheromone tables and to provide secondary guidance for ant agents proactively searching for new paths or updating the estimated quality of the paths already in use. The overall effect is a *mesh of datagram paths*, each with an associated quality estimate (the pheromone value). They are used to route data concurrently over *multiple paths* according to a *stochastic routing policy* which makes use of the pheromone values to spread data along the paths proportionally to their estimated quality.

AntHocNet has been evaluated in an extensive set of simulation tests. In particular, its behavior was studied under different conditions for network size, connectivity and change rate, radio channel capacity, data traffic patterns, and node mobility. As a benchmark, we used Ad-hoc On-demand Distance Vector routing (AODV) [16], a state-of-the-art and de facto standard algorithm. Performance was measured in

terms of data delivery ratio, end-to-end packet delay and delay jitter as *measures of effectiveness*, and routing overhead in number of control packets per successfully delivered data packet as *measure of efficiency*. AntHocNet was found to outperform AODV for all effectiveness measures on the wide range of investigated scenarios, while the efficiency of both algorithms was comparable. The good performance of AntHocNet over these heterogeneous scenarios indicate its adaptivity, while its scalability was shown by an increasing performance advantage over AODV for larger network sizes.

### 3.1.2 Topology control in sensor networks

In sensor networks *topology control* involves choosing the transmission power for each node to obtain certain global network properties. We have focused on the *minimum power topology problem*, where the aim is to assign transmission powers in such a way that all the nodes are connected by bidirectional links and the total power consumption over the network is minimized. We approached this problem according to a strategy based on *reaction-diffusion* mechanisms, widely used in mathematical models of biological systems (e.g., termite nest building [2]). In reaction-diffusion systems, elements react to specific events or situations by adapting their state and diffusing information into the environment. The presence of diffused information in turn primes reactions in other elements, leading to a sustained self-organized activity.

We have proposed the *Local Minimum Power Topology (LMPT)* algorithm [15], consisting of a reaction-diffusion mechanism that embeds an optimization phase. Each node incrementally increases its power in order to reach a minimum number of neighbors. Then, it asks its neighbors about their neighbor list and the power they use to reach them. When the node has all this information, it solves a local optimization problem involving itself and these nodes using a static optimization method [14] we developed. In the meantime the nodes in its neighborhood wait for the optimization to be concluded. At this point, according to the solution of the optimization, the node distributes the new neighbor lists and the new transmission powers for its (current) neighbors. Once they receive this information they update their state and lists, and in turn follow the same procedure, giving rise to a distributed and adaptive optimization process.

The approach has been validated using standard sensor network benchmarks. We ran tests considering networks with 1600 nodes and three different classes of node placements. We compared to the state-of-the-art algorithm *Local Minimum Link Degree (LMLD)* [8] and considered three performance indicators: total transmission power, average number of neighbors and maximum number of neighbors. In all the experiments our protocol brings a substantial gain over LMLD for all the considered indicators.

## 3.2 Algorithms for peer-to-peer networks

### 3.2.1 Data aggregation

*Aggregation* denotes the *distributed calculation of network-wide values* starting from initial values held at the nodes and

ending when, according to characteristics of the process, few or all nodes possess the final aggregate value. Such a way of proceeding allows local access to global information. Possible aggregation functions include finding extremal values, computing averages and sums, etc. This can be used to obtain information such as network size, total free storage, maximum load, etc. We developed a robust and adaptive protocol for calculating aggregates *proactively* [12]. The protocol continuously provides the value of the aggregate function to *all nodes* in an *adaptive* fashion: if the aggregate values change due to changes in the network the protocol reactivates the calculation until convergence is reached.

Our basic aggregation protocol is based on the mechanisms underlying *epidemic spreading of contagious diseases*: because of physical proximity the internal state of the host (the virus) can be transmitted to neighbors, infecting them. The iteration of the process over time, combined with dynamic neighbor relationships (e.g., due to mobility of individuals), can spread the disease across the whole population. The protocol works as follows. Each node  $p$  executes two different threads. The *active* thread periodically initiates an *information exchange* with a peer node  $q$ , selected randomly among its neighbors in the overlay network, by sending  $q$  a message containing its local state  $s_p$  and waiting for a response containing the neighbor's state  $s_q$ . That is, the information exchange is always performed in symmetric fashion. The *passive* thread waits for messages sent by another node and replies with its local state. On receiving state information from a neighbor, a node updates its state based on the value of both its current state and the received one, and according to a local aggregation function that depends on the specific global function to be calculated. We have identified which choices of the local function allow the method to calculate a number of useful aggregates, like maximum, minimum, geometric and harmonic means, network size, etc.

Through theoretical and experimental analysis we have demonstrated several desirable properties of the protocol. For instance, in the case of averaging, we proved that, if the selected peer is a globally random sample, then the variance of the approximation decreases exponentially fast, independently of network size. We also proved that the protocol is not sensitive to node crashes, while link failures were shown to only slightly slow down convergence. On the empirical side, results showed that the protocol is extremely efficient on topologies with small diameter, and confirmed the high robustness with respect to link failures, while the effect of message losses is more severe.

### 3.2.2 Content search

Search for content is a core function in P2P networks. From a content point of view, one can distinguish between *structured overlays*, in which the topology is continually adapted as a function of the content, and *unstructured overlays*, which have a random topology without clear correlation with the content and therefore require less maintenance. Structured overlays are more scalable and efficient, but less flexible and robust. Here we report work on search in unstructured overlays (see [12] for work on structured networks). Search-

ing for content in unstructured overlays is usually based on restricted flooding or random walks making use of caching and/or statistical information maintained at the nodes.

We took inspiration from the *immune system* to design efficient search heuristics based on random walk [7]. In the immune system, B lymphocytes, in response to exposure to a pathogenic agent (antigen), undergo *proliferation* and differentiation. The B cells change into plasma cells, which in turn secrete large quantities of *antibodies*, antigen-specific protein molecules finely tuned to bind to the antigen and render it harmless. *Mutation* processes further contribute to the system's effectiveness by increasing diversity in the antibody population. In our search system, called *ImmuneSearch*, the query message is seen as an abstraction of an antibody which is generated by the node initiating a search, whereas antigens are the searched items hosted in the nodes. When a search is initiated by a peer, the query message follows a random walk through the network until the content is found. At each node, the content of the query is checked versus the information profile of the node and an affinity measure is calculated. If the affinity is over a threshold, the query locally proliferates into a number of copies. A mutation function is applied before they are forwarded to selected neighbors. Proliferation has the effect of intensifying the search in the current neighborhood proportionally to the affinity, while mutation allows to look also for items closely related to the query. The effectiveness of this strategy depends on the assumption that peers with similar information are neighbors in the overlay. For this aim, the algorithm constantly *reshapes the topology of the overlay*: using the results of the affinity calculations, nodes change their position in the overlay to get closer to nodes hosting related content.

ImmuneSearch has been compared to state-of-the-art algorithms based on multiple random walks [13]. We considered a number of network scenarios (up to 10000 nodes), and performance metrics. All experimental results pointed out the superiority of ImmuneSearch.

## References

- [1] I.F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci. A survey on sensor networks. *IEEE Communications Magazine*, 40(8), August 2002.
- [2] S. Camazine, J.-L. Deneubourg, N. R. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau. *Self-Organization in Biological Systems*. Princeton University Press, 2001.
- [3] G.A. Di Caro, F. Ducatelle, and L.M. Gambardella. AntHocNet: an adaptive nature-inspired algorithm for routing in mobile ad hoc networks. *European Transactions on Telecommunications*, 16(5), 2005.
- [4] M. Dorigo, G.A. Di Caro, and L. M. Gambardella. Ant algorithms for discrete optimization. *Artificial Life*, 5(2):137–172, 1999.
- [5] F. Ducatelle, G.A. Di Caro, and L.M. Gambardella. Using ant agents to combine reactive and proactive strategies for routing in mobile ad hoc networks. *Int. J. of Comp. Intelligence and Applications*, 2005. To appear.
- [6] D. Fogel. *Evolutionary Computation*. IEEE Press, 1995.
- [7] N. Ganguly, G. Canright, and A. Deutsch. Developing Efficient Search Algorithms for P2P Networks Using Proliferation and Mutation. In *3<sup>rd</sup> Int. Conf. on Artificial Immune Systems*, September 2004.
- [8] I. Glauche, W. Krause, R. Sollacher, and M. Greiner. Continuum percolation of wireless ad hoc communication networks. *Physica A*, 325:577–600, 2003.
- [9] P. E. Heegaard, O. Wittner, and B. Helvik. Self-managed virtual path management in dynamic networks. In *Self-Star Properties in Complex Information Systems*, volume 3460 of *LNCS*. Springer-Verlag, 2005.
- [10] M. Jelasity and O. Babaoglu. T-Man: Gossip-based Overlay Topology Management. In *Proc. of the 3<sup>rd</sup> Int. Work. on Eng. Self-Org. Appls. (ESOA)*, 2005.
- [11] M. Jelasity, R. Guerraoui, A. Kermarrec, and M. van Steen. The peer sampling service: Experimental evaluation of unstructured gossip-based implementations. In Hans-Arno Jacobsen, editor, *Middleware 2004*, volume 3231 of *LNCS*. Springer-Verlag, 2004.
- [12] M. Jelasity and A. Montresor. Epidemic-Style Proactive Aggregation in Large Overlay Networks. In *Proc. of the 24<sup>th</sup> Int. Conf. on Distr. Comp. Sys. (ICDCS)*, 2004.
- [13] Q. Lv, P. Cao, E. Cohen, and S. Shenker. Search and replication in unstructured peer-to-peer networks. In *Proc. of the 16<sup>th</sup> ACM Int. Conf. on Supercomp.*, 2002.
- [14] R. Montemanni and L.M. Gambardella. Exact algorithms for the minimum power symmetric connectivity problem in wireless networks. *Computers and Operations Research*, 32(11):2891–2904, November 2005.
- [15] R. Montemanni, L.M. Gambardella, and A.K. Das. Models and algorithms for the MPSCP: an overview. In *Handbook on Theoretical and Algorithmic Aspects of Sensor, Ad Hoc Wireless, and Peer-to-Peer Networks* (J. Wu ed.). CRC Press, 2005. To appear.
- [16] C.E. Perkins and E.M. Royer. Ad-hoc on-demand distance vector routing. In *Proc. of the 2<sup>nd</sup> IEEE Work. on Mobile Computing Systems and Applications*, 1999.
- [17] E.M. Royer and C.-K. Toh. A review of current routing protocols for ad hoc mobile wireless networks. *IEEE Personal Communications*, 1999.

## Kontakt

Gianni A. Di Caro, Frederick Ducatelle, Luca M. Gambardella  
 Dalle Molle Institute for Artificial Intelligence (IDSIA)  
 Galleria 2, 6928 Manno-Lugano, Switzerland  
 e-mail: {gianni,frederick,luca}@idsia.ch  
 www: <http://www.idsia.ch>



From left to right: G. A. Di Caro, F. Ducatelle, L. M. Gambardella. The authors work at IDSIA, focusing on the application of artificial intelligence and biologically-inspired techniques to telecommunication networks and combinatorial optimization problems. L. M. Gambardella is IDSIA's director, G. A. Di Caro is post-doctoral researcher and F. Ducatelle Ph.D. student.