A Biochemical Metaphor for Developing Eternally Adaptive Service Ecosystems

Mirko Viroli, Matteo Casadei
Sara Montagna
Alma Mater Studiorum – Università di Bologna
Via Venezia 52, 47023 Cesena (FC), Italy
mirko.viroli@unibo.it
m.casadei@unibo.it
sara.montagna@unibo.it

Franco Zambonelli
Università di Modena e Reggio Emilia
Via Amendola 2, 42100 Reggio Emilia, Italy
franco.zambonelli@unimore.it

1. TOWARDS SERVICE ECOSYSTEMS

In the near future, pervasive sensing and actuating devices will densely populate our everyday environments, and will be tightly integrated with Telecom and Internet networks—also eventually contributing to blur their distinction.

In this context of tight convergence and integration, a single innovative open software platform will have to be provided to host and orchestrate in an integrated and self-managing way the execution of general-purpose pervasive services, the organization of large masses of contextual data, and the availability of several pervasive devices. Also, such an infrastructure should take into account the increasingly diverse and demanding needs of users (which will also seamlessly act as consumers and producers of data and services), and must be able to flexibly tolerate evolutions over time without requiring significant re-engineering to incorporate innovations and changing needs.

Recently, a great deal of research activity has been devoted to produce solutions to match the emerging characteristics of future networks and to solve problems related to, e.g., increasing dependability, reducing management efforts via self-* features, enforcing context-awareness and adaptability, tolerating evolution over time and eventually ensure that the overall service framework (if not all services within) can be highly adaptive and very long-lasting, even in the absence of explicit management actions. Unfortunately, most of the solutions so far are proposed in terms of "add-on", one-of solutions to be integrated in existing frameworks. The result of this process is often an increased complexity of current frameworks and the emergence of contrasting trade-off between different solutions.

In our opinion, there is need for tackling all the above problems by reformulating their foundation: we should no longer see services as strictly localised "loci" of data and functionalities, and the middleware as a provider of orthogonal services (for discovery, routing, data and context support), and where self-adaptability and evolvability are enforced via the introduction of autonomic managers. Rather, we should start taking inspiration from natural systems, where adaptability and eternal evolvability are intrinsic, are there because of the basic "rules of the game" regulating the behaviour of the overall system. Accordingly we envision an *ecosystem* model where services, data items, knowledge, and devices, are all provided with a unified view of *individuals* living in a world of ecosystem services and subject to its *eco-laws*.

The key difference among the possible approaches that can be undertaken towards the realisation of this idea stands in the metaphor adopted to model the ecosystem, its individuals, and its laws. In particular—without excluding the existence of other useful natural metaphors or the possibility of conceiving interesting non-natural ones—the main metaphors that can be adopted and have been suggested for distributed systems in general are: physical metaphors [4], chemical metaphors [2], biological metaphors [1], and ant-colony metaphors [3].

SAC'09 March 8-12, 2009, Honolulu, Hawaii, U.S.A.

2. BIOCHEMICAL METAPHOR

Developing on previous works, we here discuss how a metaphor of biochemistry (combining basic aspects of chemistry and cell biology) can be suitably exploited to address the development of service ecosystems.

Individuals as chemical reactants — The concept of chemical reactant is associated with that of an individual: the molecule kind m is the individual kind, while the concentration c is a numerical value representing the activity level of the individual—the higher it is, the more likely it will interact with others, and dually, it will become inert as activity level fades. Accordingly, individuals can be injected in the system and start interacting with others, and according to such interactions they can change shape, be continuosly generated/sustanied, or decay, as discussed in the following.

Locations as compartments — A biochemical system is typically composed of several compartments, each delimited by a membrane which filters and regulates how/whether chemical substances can cross it, either from inside to outside or viceversa. The concept of compartment can hence be associated with that of "location", i.e., an execution context for ecosystem services—a network host typically, but also wither a whole subnetwork or a part of an host. A main consequence is that individuals may be located in one location (e.g. a contextual information is bound to a network host), or may be spread in the whole network (e.g. a communication service available from everywhere, having a proxy in each node).

Eco-laws as biochemical laws — Chemical rules can be associated with eco-laws, and hence be used to make individuals change shape, decay, be continuosly generated, interact with others, and spatially diffuse in the system. An example rule is the one responsible for services getting contextualised: when the individual representing a context-dependent service reaches a certain location, it will match the individual representing contextual information, producing a "contextualised" service version. The rule achiving this would be of the kind:

$$\texttt{Service} + \texttt{Context} \xrightarrow{r} \texttt{ContextualisedService} + \texttt{Context}$$

where rate r should likely be sufficiently high so that the service gets fully contextualised pretty soon as it enters the location. An example rule used to make a data diffusing in the entire network, exploiting a sort of *toneighbour* biomechanical action, is instead as follows:

$$\texttt{GlobalData} \xrightarrow{r} \texttt{GlobalData} + toneighbour(\texttt{GlobalData})$$

by which at rate r `GlobalData` creates a clone molecule `GlobalData` and sends it outside the membrane towards a neighbouring compartment, namely, the data gets diffused to the neighbourhood on a step-by-step basis.

3. BASIC PATTERNS

Based on the above metaphor a brief set of simple examples of eco-laws can be introduced.

Production/decay equilibrium — Concentration can be understood as an activity level, used for certain individuals to achieve self-organisation in conjunction with the typical mechanisms of positive and negative feedbacks. Suppose a data item is to be deployed in a certain location, and its pertinency should depend on both time and exploitation, namely, its activity should both decrease as time passes and raise as the data is successfully exploited. The two mechanisms should be balanced, e.g., by the combination of the following rules:

$$\texttt{Service} \xrightarrow{r\_decay} 0$$
$$\texttt{ServiceFeedback} \xrightarrow{r\_feed} \texttt{Service} + \texttt{ServiceFeedback}$$

Execution runs—or simulations—would show that the activity level of `Service` reaches an equilibrium around a given value (the ratio between r_feed and r_decay); when the `ServiceFeedback` individual is dropped `Service` level fades to 0 with negative exponential trajectory.

Self-adapting to servive requests — Suppose now that while the service activity level is in equilibrium, a huge number of individuals representing requests for exploiting the service are injected in the location. In that case it would be interesting to temporarily increase the service activity level so as to ensure that requests are more quickly addressed: as soon as requests are served, the activity level can return back to the standard level. We add the following law:

$$\texttt{Service} + \texttt{Request} \xrightarrow{r\_use} \texttt{Service} + \texttt{Service} + \texttt{Reply}$$

The corresponding behaviour is such that as a burst of requests is injected into the system the service activity level immeditalely increases, but then it goes back to standard activity level as soon as all such requests have been served.

Field-like diffusion — Not all services are to be placed into a single location, but rather, many of them should spread within a network subpart, that is, they should diffuse from a source locality into the neighbourhood. This mechanism is rather well known in the self-organisation area as "field diffusion" [4], resembling the way in which fields of physics are used to create a spatially distributed interaction medium for particles. A similar mechanism is also useful in biology: proteins and other substances disseminate into a tissue of cells on a step-by-step basis, and similarly, fields of signals are created in early stages of embryogenesis to properly move cells as to shape the life form. In service ecosystems this mechanisms is key to obtain globally perceived services. Diffusion can be supported by a rule like the one previously seen:

$$\texttt{Service} \xrightarrow{r\_move} \texttt{Service} + toneighbour(\texttt{Service})$$

By this rule, starting from a location hosting `Service`, a persistent field gets pumped in the network, decreasing as the distance of the pumping host grows.

4. STATE OF THE RESEARCH

Current research efforts are devoted to the development of a computational model, of a high-performance simulator for experimenting with the biochemical metaphor, and to the identification of new practical biochemical patterns for ecosystem services. Future works will be devoted to implementing an actual infrastructure for running system ecoservices, based on the architectural approach described in [5].

5. REFERENCES

[1] O. Babaoglu, G. Canright, A. Deutsch, G. A. D. Caro, F. Ducatelle, L. M. Gambardella, N. Ganguly, M. Jelasity, R. Montemanni, A. Montresor, and T. Urnes. Design patterns from biology for distributed computing. *ACM Trans. Auton. Adapt. Syst.*, 1(1):26–66, 2006.

[2] A. P. Barros and M. Dumas. The rise of web service ecosystems. *IT Professional*, 8(5):31–37, 2006.

[3] M. Casadei and M. Viroli. Applying self-organizing coordination to emergent tuple organization in distributed networks. In *Second IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO'08)*, pages 213–222, Venice, Italy, 20–24 Oct. 2008. IEEE Computer Society.

[4] M. Mamei and F. Zambonelli. *Field-based Coordination for Pervasive Multiagent Systems*. Springer Verlag, 2006.

[5] F. Zambonelli and M. Viroli. Architecture and metaphors for eternally adaptive service ecosystems. In *Intelligent Distributed Computing, Systems and Applications*, volume 162/2008 of *Studies in Computational Intelligence*, pages 23–32. Springer, 2008.