Self-Management and the Many Facets of "Non-Self"

Franco Zambonelli
Dipartimento di Scienze e Metodi dell'Ingegneria
Università di Modena e Reggio Emilia
franco.zambonelli@unimore.it

The difficulties in dealing with increasingly complex information systems that have to operate in dynamic operational environments calls for self-management properties or, more in general, for the integration of "self-*" features (e.g., self-configuration, self-adaptation, self-healing) in software and information systems.

The common perspective of nearly all "self-*" approaches is that of considering human beings as "non-self" from the information system perspective (Figure 1-a). Indeed, one of the very common goals of all approaches is to move humans out-of-the-loop, by making information systems able to perform in an autonomous way all that kind of (costly and often too complex to bear) human activity related to configuring and maintaining information systems so as to have them properly working under all conditions. However, beside this common perspective, a number of diverse conceptions co-exist for what should be actually considered "self" and what, beside humans, "non-self".

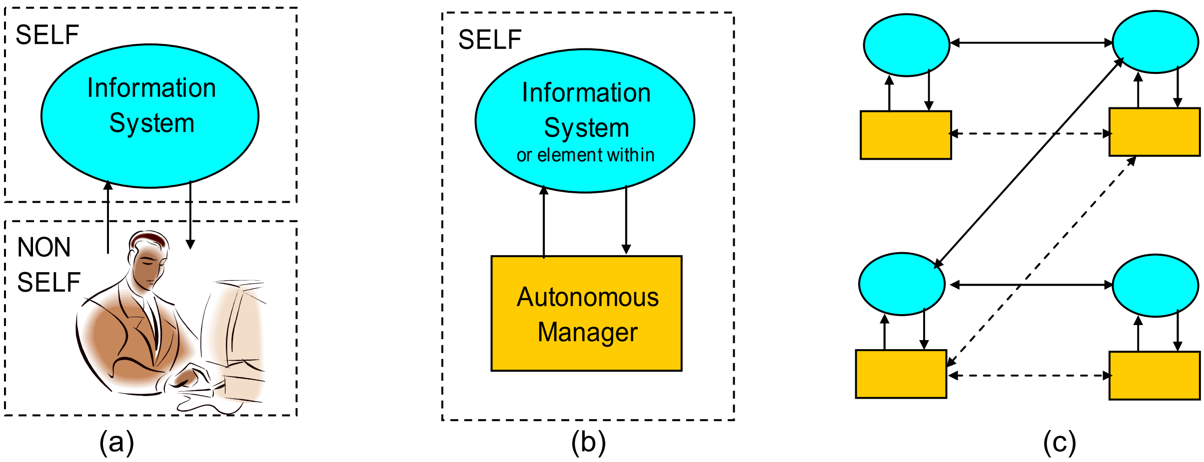

Figure 1. Moving Humans Out-of-the-Loop. (a) Humans as managers of information systems are considered "non-self" from the self-management perspective. (b) The autonomic computing perspective considers substituting "non-self" human managers with "self" surrogates in the form of digital autonomous managers. (c) Autonomous managers can be multiple and distributed, and can interact with each other to enforce distributed self-management activities.

The Autonomic Computing Perspective

The more applied industrial perspective – as reflected, e.g., by the Autonomic Computing initiative at IBM [2] – tends to focus on the basic viewpoint of considering humans as the only "non-self" beings. The overall vision is that humans should be substituted by digital surrogates, able to perform monitoring, configuration, and maintenance activities in autonomy and without human intervention.

At the level of individual information systems (or of individual components within them), this occurs via proper components that monitor what's happening, and can autonomously plan actions to re-configure and restore the system as needed, in a sort of continuous control loop (Figure 1-b). At the level of distributed computing systems, this can occur via a set of distributed managers associated to different distributed elements that, by exchanging information with each other and by orchestrating their respective actions, can ensure specific functional and non-functional properties in the overall behaviour of the distributed system (Figure 1-c).

Such an approach leads to a conceptually very clean architecture for self-managing systems, well grounded on the past lessons of operating systems and of distributed systems researches. Coupling traditional approaches to monitoring and resource management with artificial intelligence techniques for planning and knowledge-management as well as with multiagent techniques for automated negotiation may lead, in the short term, to the actual release of seemingly self-managing systems. However, an architecture based on autonomous managers that are logically separated from the components they control introduces several potential drawbacks. In fact, accounting for all possible contingencies and being able to react to them with proper actions to ensure continuous functioning may end up in a heavy-weight architecture and/or in slow and inappropriate reactions, undermining at the very basis the self-management capabilities of the system.

The Self-Organization Perspective

To some extent, the limitations of the autonomic computing perspective mainly derive from the fact that it inherits the basic architecture of the traditional human-based management approaches. Even if humans are no longer in the loop, one may be tempted to say that the autonomous managers of Figures 1-b and 1-c are digital "non-self", being aliens to the information system itself.

Self-organization approaches to self-management (exemplified by the research articles of this issue) start from this viewpoint and consider that a system should be able to self-manage by its own very nature, and not by external intervention of "non-self" entities, even if digital ones. To this purpose, they take inspiration from natural adaptive systems and from their capabilities of self-organizing their global activities into highly adaptive functional patterns. Systems such as bacterial colonies, insect colonies, embryos and organs, exhibit global functional patterns of activities – autonomously emerging from simple local activity rules and local inter-components interactions – and are able to react to both internal (e.g., death of components) and external (e.g., environmental perturbations) contingencies by re-shaping the system so as to ensure preservation of the global functional pattern. Since several of these natural phenomena find a natural mapping to some functional problems in modern and distributed information systems [1], the result is in a robust self-managing information systems in which the functional parts and the management parts are seamlessly integrated into the same components (Figure 2-a).

The key advantage of exploiting self-organization for the building of self-managing systems is that the overall architecture of the system is simpler and more light-weighted than in autonomic computing approaches. Also, since self-management properties are intrinsic in the nature of systems, no complex planning or knowledge management activities are required to properly react to specific or unforeseen contingencies. However, despite the increasing number of success stories, self-management via self-organization is not a panacea, and current self-organization approaches still suffer from several limitations. First, the rich catalogue of natural phenomena with useful application in information systems does not eradicate the "solutions in search of a problem" nature of the engineering process: the reverse engineering of some natural phenomena that incidentally happens to map into some kinds useful distributed application. General methodologies for building by "direct engineering" a self-organizing system that solves specific problems are still missing. Second, most of current approaches to self-organization can enforce a single or a limited set of self-managing functionalities, i.e., those directly related the self-organizing functional pattern, but fail in properly accounting the diverse functionalities and possibly competing needs that may be present in complex real-world information systems.

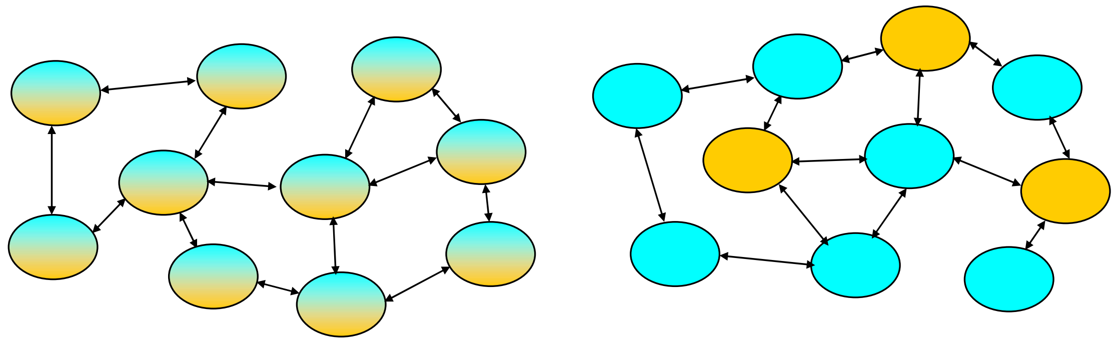

(a) (b)

Figure 2. Self-organization vs. Ecological Approaches to Self-Management. (a) In a self-organization approach, a uniform set of self-organizing components locally interact with each other and act both a functional and management components. (b) In an ecological approach, self-organizing components live and interact with each other and with additional "manager" components that may somehow control and direct the overall behaviour of the system.

The Ecological Perspective

The current limitations of self-organization approaches may suggest that most real-world self-management information systems still require the presence of external "managers", able to acquire a more global view of the whole system and of all its needs. However, an alternative approach could be envisioned to get the best of both autonomic and self-organization approaches. To avoid the introduction of complex and heavyweight managers without missing control opportunities over a system and, at the same time, to preserve the cleanness of self-organizing architectures, a possible solution is to smooth the "non-selfishness" of autonomous managers by making them become first-class citizen in self-organizing systems.

Basically, the idea is to inject in a self-organizing system additional "manager" components that, while living inside the system and interacting with other components as if they were autochthonic components, could have a differentiated behaviour, enabling them to somehow affect the global behaviour of the system itself (Figure 2-b). The idea is that these components should not undermine the basic self-organizing (and thus self-managing) nature of the systems, and should rather be part of it. However, by exploiting some more knowledge and abilities than those available to normal components, they could: (i) direct the evolution of the self-organizing system from any of several possible configurations to specific desirable configurations (i.e., those better accommodating needs that the self-organizing system per se is not able to account for); and/or (ii) improving its effectiveness in self-reorganizing upon contingencies.

In our research group we have conducted some preliminary experiments in this direction. On the one hand, we have shown how it is possible to globally direct the dynamical evolution of cellular automata by simply modifying a very limited percentage of the cells [3]. On the other hand, we have shown how, in field-based coordinated systems, it is possible to escape from sub-optimal configurations by having some of the components capable of "reasoning" about the local shape of fields, other than simply reacting to them [4]. Other interesting experiences are being performed at the Laboratory of Theoretical Biology of the Université Libre de Brussels, showing how a few robotic ants in a colony of real ants can affect the behaviour of the whole colony.

Beyond the horizon, one could imagine a scenario in which our networks will be like large ecosystems, and will host multiple "specimens" of complex self-organizing systems, co-existing over the same resources and interfering with each other in hardly predictable ways. Since these systems will be decentralized, without clearly identifiable stakeholders, and eternally running, the only solution to enforce some forms of control over them, and to have the self-management features of each individual system properly co-exist with more global forms of self-management, will be that of populating the ecosystem with additional specimens of "manager" components. After all, such an approach is already at work in agriculture, where the spread of parasites is contained with the introduction of natural predators, as well as in marketing, where it is common practice recruiting "opinion leaders" among normal people to have them promote specific products from within tribes. In any case, for such an ecological approach to become a usable practice for the management of complex information system, a long way of research experiences and production of suitable engineering tools is needed.

An interesting consequence of the ecological perspective is that it is likely to undermine the basic initial assumption of human "non-selfishness". As information systems are more and more pervasive and integrated with both the physical and the social worlds, humans will have to be necessarily considered an integral specimen of the information system ecology, and their activities and behaviours will directly affect the overall behaviour of the system and its self-management properties. In other words, humans will get back into the management loop as first-class "self" entities, even if implicitly and without direct management responsibilities. As a simple example, consider students accessing the Wi-Fi mesh of a campus with their laptops. While the mesh self-organizes its activities and re-distributed connections to provide services of suitable quality to everybody, most likely some of the students will try to optimize their own positions in the campus so as to get better connectivity. To some extent, one could say that these students and the system implicitly cooperate for the optimal self-management of the network. An analysis of the implications of these aspects, though, would require much more room than the few pages of this article and much more inter-disciplinary competences than I actually have.

As un-concluding remarks, I can only say that the spectrum of possible perspectives on self-management makes it very hard to predict what the future will be. If I had to, I would bet on (i) autonomic computing approaches to prevail in the short term, (ii) being gradually integrated with self-organizing approaches in the medium term and (iii) eventually subsumed by ecological approaches in the long term. Whatever the case, there's plenty of room for exciting researches along all these directions.

References

[1] O. Babaoglu, G. Canright, A. Deutsch, G. Di Caro, F. Ducatelle, L. Gambardella, N. Ganguly, M. Jelasity, R. Montemanni, A. Montresor, "Design Patterns from Biology for Distributed Computing", in *Proceedings of the European Conference on Complex Systems*, Paris (F),

November 2005.

[2] J. Kephart, D. M. Chess, “The Vision of Autonomic Computing”, *IEEE Computer*, 36(1):41-50, January 2003.

[3] M. Mamei, A. Roli, F. Zambonelli, “Emergence and Control of Macro Spatial Structures in Perturbed Cellular Automata, and Implications for Pervasive Computing Systems”, *IEEE Transactions on Systems, Man, and Cybernetics – Part A*, 36(5):337-348, May 2005.

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