The Ecology of the Grid

Fabrice Saffre
British Telecommunications plc.
Pervasive ICT Research Centre
fabrice.saffre@bt.com

José Halloy and Jean Louis Deneubourg
Université Libre de Bruxelles
Service of Social Ecology
jhalloy, jldeneub@ulb.ac.be

Abstract

In this paper, we show how self- mechanisms give rise to complex but predictable and therefore steerable global system behavior in a co-operative computing environment. We simulate the operation of and interactions between a set of networked autonomic devices. These are used as access points to a number of services, have the ability to accept or delegate execution of the associated tasks, and can adjust their internal state in response to the demand. We study the emergence of co-operation, and find that it spontaneously occurs when specific conditions are met that allow individual devices to focus on performing a single task, sacrificing their ability to efficiently perform others.*

1. Introduction

Because it only relies on basic scripts, it has often been wrongly assumed that biological self-organization is only found in systems comprised of extremely simple units. In reality, it often combines with (or even makes use of) advanced individual information processing capabilities to produce higher levels of collective complexity [1]. Yet when using Nature as a blueprint for artificial systems, and even when deliberately ignoring the sophistication of individual behavior, understanding how simple rules of interaction translate into system properties can already be a challenge.

Fortunately, on an abstract level of description, one finds that a surprisingly small number of such simple rules govern nearly all biological systems (from the cellular level to animal societies) and produce efficient emergent collective patterns for resource management, task allocation, differentiation, synchronisation, clustering etc. These generic rules effectively act as "building blocks", the combination of which presiding to a wide variety of self-organizing system behaviors.

Our objective is to demonstrate that the same paradigm can successfully be applied to study, understand, anticipate and ultimately control the dynamics of autonomic systems. Indeed, the fundamental similarity between a community of living organisms and a network of devices individually governed by autonomic principles is obvious: both are aggregations of agents trying to maximise their performance, in an environment where efficiency is largely determined by the nature of their interactions with other agents.

We chose grid computing [2], in the widest possible sense of co-operative resource sharing, as the generic background of our demonstration primarily because of its long-lasting association with biological metaphors [3].

2. Algorithm and simulation results

We are considering the case of a group of devices acting as access points for a set of services, each of which corresponding to a single task. Every device is capable of interacting with its four first neighbours in a square lattice (periodic boundary conditions), of adjusting its own configuration to modify its task-specific performance, and of deciding whether or not to perform a task when requested to provide the corresponding service by agents external to the system (e.g. users). Providing the requested service is mandatory, so a device can only refuse to perform a task if it is able to successfully delegate execution to one of its neighbours.

The decision of accepting or refusing to perform a task is governed by a "local reasoning loop" inspired from previous work [4]. Basically, the probability $P_{i,j}$ that device i accepts a request for task j is simply given by $P_{i,j} = 1/(1+(r_{i,j}/r_c)^{\alpha})$, where $r_{i,j}$ is the Euclidian distance between the location of i and that of the optimal configuration for performing j in an abstract, two-dimensional task space (r_c and α are parameters).

In our model, every time that a device performs a task, it re-configures itself, "moving closer" to the corresponding optimum. Without co-operation, and assuming identical demand for all services, this form of specialization simply tends to bring every device in the vicinity of the point that is equidistant from the *n* task optima corresponding to the *n* services for which it is a designated access point. We call this "node-centric autonomic scenario", as it effectively means that every device individually tunes its own internal state so as to reach the best possible compromise between conflicting requirements.

However, when coupled with the ability to delegate a task to another unit and learn (via the simplest form of trial-and-error) "who's good at doing what job", the same self-configuring behavior is enough to induce specialization. We call this "co-operative autonomic scenario", as it effectively means that all participating devices start relying on each other's support, increasing overall efficiency at the expense of individual versatility. This process requires no central information brokering or task scheduling, whereby devices would be "instructed" to differentiate into specialists and tasks assigned accordingly. Instead, the system adapts itself exclusively via local, unsupervised decision-making and trial and error interactions between first neighbours.

The underlying mechanism is a simple, non-explicit positive feedback loop, which, because of the direct relationship between task-specific efficiency and acceptance probability amplifies heterogeneities present in initial conditions and/or introduced into the system by random fluctuations of the demand in the early stages of its history. In the right conditions, this leads the system as a whole to turn into a "community" of mutually dependent but highly efficient task specialists, instead of an unstructured group of "all-rounders".

It is key to understand that co-operative effects emerge from the system's ability to evolve a set of associations that are acceptable to a majority of participating devices, even though the asymmetrical nature of individual "task request" channels and the absence of any direct reciprocity may suggest a form of parasitism rather than symbiosis. The main reason is that the quality of a "helper's" work is directly proportional to its ability to differentiate, which in turn depends on its own success when trying to delegate the execution of other tasks.

In the right conditions, the result is the spontaneous formation of co-operative teams, as illustrated by the example shown on fig. 1, where the greyscale (darker = better) is an indication of link quality (in terms of the frequency of positive answers to task requests) and of device overall performance.

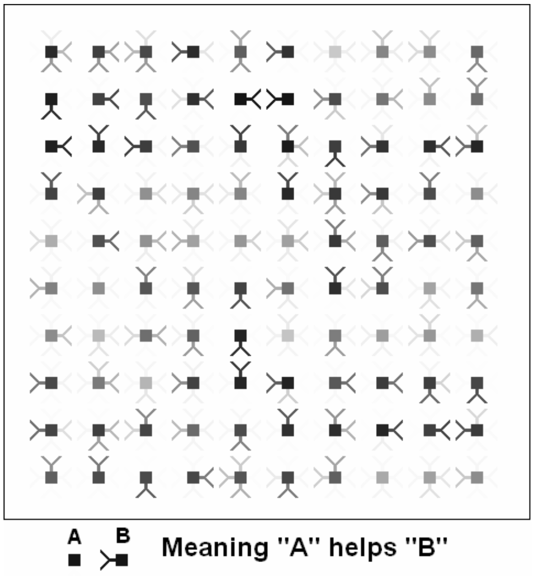

Figure 1. Snapshot of the "map" of help relationships having developed in the co-operative scenario after 2×10^6 time-steps.

Note that, typically, only devices that are doing relatively well in terms of their own ability to provide good quality of service (i.e. through a combination of specialization and successful delegation) are themselves consistently used as "helpers". In contrast, access points that found themselves "isolated" due to circumstances spontaneously revert to the "node-centric autonomic behavior", achieving lower (but better than random) overall efficiency by simply adapting their configuration to the local demand.

3. References

[1] Camazine, S., Deneubourg, J.L., Franks, N., Sneyd, J., Bonabeau, E. and Theraulaz, G., "*Self-Organization in Biological Systems*", Princeton University Press, 2001.

[2] Foster, I., Kesselman, C. and Tuecke, S., "The Anatomy of the Grid", *International Journal of Supercomputer Applications*, 15(3), 2001.

[3] Montresor, A., Meling, H. and Babaoglu, O., "Messor: Load-Balancing through a Swarm of Autonomous Agents", *Proceedings of the 1st Workshop on Agent and Peer-to-Peer Systems*, 2002.

[4] Saffre, F. and Blok, H.R., "SelfService – A Theoretical Protocol for Autonomic Distribution of Services in P2P Communities", *Proceedings of the 1st International Conference on Autonomic Computing*, 2004.