Chapter LI
Virtual Organization Support through Electronic Institutions and Normative Multi-Agent Systems

Henrique Lopes Cardoso
University of Porto, Portugal

Ana Paula Rocha
University of Porto, Portugal

Eugénio Oliveira
University of Porto, Portugal

ABSTRACT

The multi-agent system (MAS) paradigm has become a prominent approach in distributed artificial intelligence. Many real-world applications of MAS require ensuring cooperative outcomes in scenarios populated with self-interested agents. Following this concern, a strong research emphasis has been given recently to normative MAS. A major application area of MAS technology is e-business automation, including the establishment and operation of business relationships and the formation of virtual organizations (VOs). One of the key factors influencing the adoption of agent-based approaches in real-world business scenarios is trust. The concept of an electronic institution (EI) has been proposed as a means to provide a regulated and trustable environment, by enforcing norms of behavior and by providing specific services for smooth inter-operability. This chapter exposes our work towards the development of an agent-based EI providing a virtual normative environment that assists and regulates the creation and operation of VOs through contract-related services. It includes a presentation of the EI framework, knowledge representation structures for norms in contracts, and a description of two main institutional services, namely negotiation mediation and contract monitoring.

INTRODUCTION

The multi-agent system (MAS) paradigm has become a prominent approach in distributed artificial intelligence. These systems start with the individual—the agent—and evolve to populated environments where the most important feature is the interaction between agents.

The concept of intelligent agents is today dominant in artificial intelligence (Russell & Norvig, 2003). Agents are described as entities having a set of intrinsic capabilities, such as *autonomy* (control of own decision making), *reactivity* (response to changes in the environment), and *pro-activeness* (goal-directed behavior) (Wooldridge & Jennings, 1995). Moreover, the so-called strong notion of agency considers agents as entities having mental states, including *beliefs, desires,* and *intentions* (the BDI architecture; Rao & Georgeff, 1995). Another important capability—*social ability*—emphasizes the fact that agents do not act in an isolated environment, and will inevitably have to interact with other agents.

Thus, while agent theory has its inspirations in psychology and cognitive science, MAS research is influenced by organizational and social sciences, distributed computing, and economics. Furthermore, many MAS researchers look for inspiration in nature, both from collective cooperative behavior in groups of animals (e.g., ant colonies) and from our social interactions as (not necessarily cooperative) human beings.

Typical MAS applications include inherently distributed or complex domains. While some problems require system architectures including cooperative agents developed so as to accomplish an overall goal, in other cases agents may represent independent self-interested entities, with no presupposed cooperation besides mere interaction efforts. The former types of problems may be addressed through a centralized design, producing a top-down specification of a MAS environment with an overall purpose. The latter types are usually conceived as open environments, where heterogeneous agents arising from different sources interact either cooperatively or competitively. In this setting, agents may form organizations that dynamically emerge in a bottom-up fashion from the individuals, which together agree, usually through a negotiation process, to cooperatively perform some task not doable individually.

As one might expect, although decentralized and dynamic systems are much more appealing, they must be handled with hybrid approaches, since a minimum set of requirements is necessary to allow for heterogeneous and independently developed agents to successfully interact. One way of achieving such a common milieu is by defining communication standards, such as those proposed by FIPA (2002). However, an important issue arises when attempting to apply agents in real-world settings: how to ensure cooperative outcomes in scenarios populated with self-interested agents. A possible answer to this problem is to regulate the environment, enforcing appropriate types of agent behavior. This should provide a level of trust necessary for the development of real-world applications of open MAS.

Following these concerns, a strong research emphasis has been given recently to normative MAS (Boella, van der Torre, & Verhagen, 2005). A normative system is a set of interacting agents whose behavior can usefully be regarded as governed by norms (Jones & Sergot, 1993). Agents are subject to these norms, which influence their decision making. Therefore, besides their goals, agents must take into account the norms that apply to them. However, considering autonomy as a central property of agents, norms are often used as a means to regulate the environment by providing incentives for cooperative behavior through norma-