Normative Run-Time Reasoning for Institutionally-Situated BDI Agents

Tina Balke[1,2], Marina De Vos[2], and Julian Padget[2]

[1] University of Surrey, Centre for Research in Social Simulation
t.balke@surrey.ac.uk
[2] University of Bath, Dept. of Computer Science
{mdv,jap}@cs.bath.ac.uk

Abstract. Institutions, also referred to as normative systems, offer a means to govern open systems, in particular open multi-agent systems. Research in logics, and subsequently tools, has led to support for the specification, verification and enactment of institutions. Most effort to date has focused on the design-time properties of institutions (either on the normative or the system level), such as whether a particular state of affairs is reachable or not from a given set of initial conditions. Such models are useful in forcing the designer to state their intentions precisely, and for testing (design-time) properties. However, we identify two problems in the direct utilization of design-time models in the governance of live (run-time) systems: (i) over-specification of constraints on agent autonomy and (ii) generation of design-time model artefacts. In this paper we present a methodology to tackle these two problems and extract run-time reasoning components from a design-time model. We demonstrate how to derive an event-based run-time model of institutions that can be incorporated into the capabilities of autonomous BDI agents to address the issues above in order to realize practical norm-governed multi-agent systems.

1 Introduction

The motivation for this work is two-fold: the first is the goal of the run-time governance of open distributed systems and the second is a case study of such a system: using institutions to govern the interaction of participants to demonstrate the economic viability of future mobile phone networks (called wireless grids)

The general idea in these future mobile phone networks is that mobiles dynamically construct ad-hoc wireless grids with the objective of achieving (i) faster download times by splitting content into parts, downloading some using 3G and acquiring the rest from nearby phones using wifi (ii) reducing power consumption by trading off high-cost 3G communication for low-cost wifi communication [7].

The main challenge within this context is how to encourage participants to contribute actively to the collective by downloading their share of parts via the high-cost 3G link and then to share them with the others, rather than free-ride by only receiving parts from the other participants in the system.

S. Cranefield et al. (Eds.): COIN 2011, LNCS 7254, pp. 129–148, 2012.

We propose an agent-based simulation to analyse cooperation effects in these mobile phone networks, using institutions as a mechanism to encourage and enforce cooperation. Results from the simulation can then be used to determine the viability of the system.

Typically institutions and agents co-exist in a state of tension: agents are (supposed to be) autonomous, while institutions constrain autonomy. Often, in norm-governed MAS, this tension is alleviated by regimenting agents and their actions [12], thus not allowing any norm-deviation. We, in contrast, use a more social form of institution. An agent can query institutional properties at run-time in order to examine how situations were achieved, or can possibly be achieved in the future, to determine which normative context is applicable to their current situation and to evaluate possible futures and make a decision on being norm-compliant or not, on the basis of this information.

Given the event-based nature of the simulation, the institutional approach described in [3] offers a suitably compatible model, as well as a complementary computational model that could be adapted to provide agents with information about the institutional state. At the same time, we also needed a suitable agent architecture, with a programming model that would fit the requirements for both being able to process institutional events and taking a goal-driven approach to the tasks to be fulfilled in the simulation.

Traditionally, when trying to analyse normative effects on a system, the real world is formalized as two *separate* (i.e. no dependencies) models: a system model and a normative/institutional model of which only the design-time properties are analysed *separately*. While useful, this can be problematic when wanting to analyse the interplay between agents and the institution. Furthermore, it poses the problem of how and when to account for run-time effects. Thus, in contrast to the separate analysis of the normative and system models, we are interested in an integrative and coherent analysis of the two models indvidually and jointly. Thus, we approach the institutional and system modelling in two phases:

1. *Design-Time Model:* we start traditionally with a design-time model, which in contrast to the standard approaches, integrates the system and normative model and allows for design-time reasoning about the interplay between agents on the system level and the institutions on the normative level. For instance, we build an institutional model of the wireless grid concept to evaluate whether it makes sense to pursue the idea, i.e. whether enforcement improves collaboration even in the most favourable circumstances. A design-time model is less labour intensive than implementing a full system. This model hard-codes simplifications of the environment in which the agents interact, but it can be used for validation purposes and helps to expose requirements issues.
2. *Run-Time Model:* Using the design-time model as a starting point, we then derive a run-time model. This is created by removing all but the normative information and domain facts from the design-time model, in order to avoid design artifacts and restrictions with respect to agent autonomy.