IRON: A Machine for the Automated Synthesis of Normative Systems

(Demonstration)

Javier Morales[1,2] and Maite Lopez-Sanchez[1]
[1]MAiA Department
Universitat de Barcelona
Barcelona, Spain
jmorales@iiia.csic.es
maite@maia.ub.es

Juan A. Rodriguez-Aguilar
[2]Artificial Intelligence Research Institute (IIIA)
Spanish Council of Scientific Research (CSIC)
Campus UAB. Bellaterra, Spain
jar@iiia.csic.es

Michael Wooldridge
Dept. of Computer Science
University of Oxford
Oxford, United Kingdom
mjw@cs.ox.ac.uk

Wamberto Vasconcelos
Dept. of Computing Science
University of Aberdeen
Aberdeen, UK
wvasconcelos@acm.org

ABSTRACT

The automated synthesis of norms for coordination of multi-agent systems remains an open and complex problem. In this paper we present the Intelligent Robust On-line Norm Synthesis Machine (IRON), a system whose goal is the automated synthesis of norms. IRON is capable of synthesising norms that are at the same time effective (to ensure coordination) and necessary (to avoid over-regulation). IRON has been tested on a simulated traffic scenario to successfully synthesise norms that help cars avoid collisions. IRON is equipped with visualization features that provide support for an intuitive and informed monitoring of the synthesis process.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence —*Multiagent Systems*

Keywords

Norms - Normative Systems - Norm Synthesis

1. INTRODUCTION

Since the seminal work in [6], the problem of norm synthesis (i.e., determining the set of norms that avoid conflicting states) has attracted considerable attention within the MAS community. We differentiate two strands of work tackling this problem: the *off-line* and *on-line* norm synthesis approaches. On the one hand, off-line approaches (e.g. [6]) aim at synthesising norms for a MAS that constrain the behaviour of agents while ensuring the achievement of global system goals. Nonetheless, off-line design is not appropriate to cope with open MAS, whose composition and state space change with time. On-line norm synthesis approaches (e.g. [4][5]) try to overcome such limitations by synthesising norms that regulate a MAS at run-time instead of at design time. It considers that agents collaboratively choose their own norms out of a space of possible norms. A norm is considered to have emerged when a majority of agents adopt and abide by it.

Against this background, we propose a novel system, the so-called IRON (Intelligent Robust On-line Norm Synthesis Machine),

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, Ito, Jonker, Gini, and Shehory (eds.), May, 6–10, 2013, Saint Paul, Minnesota, USA.

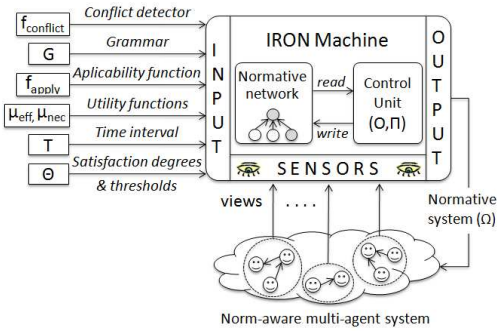

Figure 1: IRON**'s architecture.**

for the on-line synthesis of norms (demo available at [1]). IRON produces norms for the agents in a MAS that characterise necessary conditions for coordination, while avoiding over regulation. IRON synthesises norms that are both *effective* and *necessary*. IRON is also capable of generalising norms. By generalising norms and discarding unnecessary norms, IRON yields *concise* normative systems. As a result, IRON manages to successfully synthesise norms that are both effective and necessary, even in the presence of non-compliance behaviours in a MAS. The visualisation features with which IRON is equipped provide the MAS engineer with support for an intuitive and informed monitoring of the synthesis process.

2. IRON'S OPERATION

This section outlines the operation of IRON. For a full account and details of its operation, refer to [3]. Given a norm-aware multi-agent system (NA-MAS), IRON operates by continuously running the following steps: (1) it monitors the NA-MAS operation in search for conflicts; (2) it decides upon the addition of brand new norms to the current (initially empty[1]) normative system (defined as the current set of active norms that regulate the system); (3) it evaluates whether the effectiveness and necessity of the normative system are within expected thresholds; (4) if required, it refines the normative system; and (5) it makes the normative system available to agents. Notice therefore that IRON continuously searches for a normative system *on-line*, as agents in the NA-MAS operate.

IRON is based on four main components: (i) a grammar to synthesise new norms; (ii) a normative network (a data structure to

[1]The approach would also work if the normative system is initialised with a set of norms provided at design time.

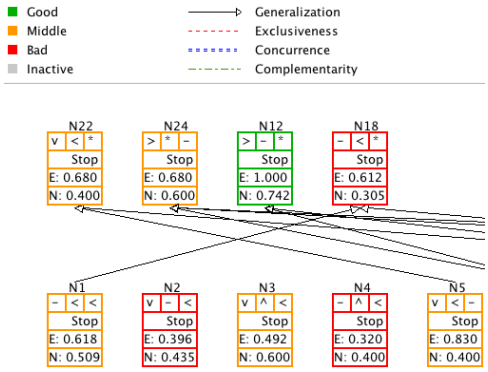

Figure 2: A normative network as displayed by IRON.

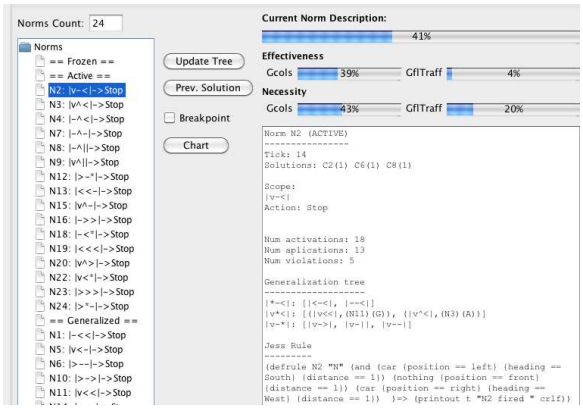

Figure 3: IRON's norm information panel.

represent normative systems); (iii) a set of operators that allow to transform a normative system into another; and (iv) a strategy that specifies when to use such operators. Figure 1 shows how these components are located in IRON's architecture, which is a refinement of the norm cycle described in [2].

IRON represents normative systems by means of a graph-based data structure, named a *normative network*, whose nodes stand for norms and whose edges stand for relationships (generalisations in this paper) between norms. Norms in a network may be either *active* or *inactive*. We consider that the active norms in a normative network represent a normative system.

The norm synthesis process starts by detecting conflicts in a MAS that is observed by IRON. For each detected conflict, the strategy synthesises a new norm in order to avoid it in the future. Subsequent norm utilities are iteratively evaluated by computing the effectiveness and necessity of each norm at each time t.

On the one hand, IRON measures the effectiveness of *applied* norms based on their outcomes. It evaluates the *cumulative* effectiveness of a norm according to the following principle: the higher the ratio of *successful applications* (applications not leading to conflicts) of a norm, the higher the effectiveness increase. On the other hand, IRON assesses the *cumulative* necessity of a norm according to the following principle: the higher the ratio of *harmful violations* (violations leading to conflicts), the more necessary the norm.

Finally, IRON's strategy performs a *normative system refinement*, which yields a new normative system by transforming the normative network via specialisations and generalisations. With this aim, the strategy keeps track of the effectiveness and necessity of the norms in the normative network during a period of time T. Then, the refinement task amounts to implementing the following rules:
(1) A norm is *specialised* (or *deactivated* if it has no children in the normative network) provided that either its effectiveness *or* necessity have not been good enough during a period T. This occurs when the effectiveness *or* necessity of some of its children have not been good enough either.
(2) A set of norms are *generalised* provided that: (i) they all relate to the same norm (parent) in the normative network; (ii) they are the possible child norms of the parent norm; (iii) their effectiveness *and* necessities have all been good enough during a period T.

IRON provides monitoring facilities, through the norm information panel (Fig. 3), to track all the details regarding each norm (evaluation, encoding, state, etc.). Furthermore, it also provides an additional monitoring facility to track the evolution of a normative network, which includes the generated norms along with their evaluations and relationships. Figure 2 shows a sample of normative network as shown by IRON.

3. DEMO SCENARIO

For this demo, we connect IRON with a MAS simulation. Thus, our demo scenario is a NA-MAS simulation of a traffic junction composed of two orthogonal roads represented by a 19×19 grid. Each road has two 19-cell lanes (one per direction). Each agent is a car that travels along the grid by following a random trajectory (i.e. random entry and exit points). The demo allows a user to configure the traffic density and the probability of norm violations prior to start a simulation. During the simulation, the user can employ the above-mentioned IRON's visualisation tools as well as IRON's charts to track the synthesis process. Furthermore, during a simulation, IRON's log facilities record the data required to analyse results in terms of convergence to and utility of normative systems.

4. CONCLUSIONS

IRON is a novel system for the on-line synthesis of norms. It synthesises norms for the agents in a MAS that guarantee conflict-less coordination while avoiding over regulation. For this purpose, IRON employs effectiveness and necessity as the measures that characterise the quality of a normative system. Furthermore, IRON is capable of generalising norms. By keeping effective norms, generalising norms, and discarding unnecessary norms, IRON yields *effective* and *concise* normative systems. Overall, IRON shows in an intuitive and comprehensive way the norm synthesis process.

5. REFERENCES

[1] IRON: a machine for the automated synthesis of normative systems. `http://www.youtube.com/watch?v=QwwQYqu4EYo&feature=youtu.be`, 2013.

[2] J. Morales, M. López-Sánchez, and M. Esteva. Using Experience to Generate New Regulations. In *IJCAI2011*, pages 307–312. AAAI Press, USA, 2011.

[3] J. Morales, M. López-Sánchez, J. A. Rodriguez-Aguilar, M. Wooldridge, and W. Vasconcelos. Automated synthesis of normative systems. In *Proceedings of the 12th AAMAS*, 2013. To be published as a full paper.

[4] B. Savarimuthu, S. Cranefield, M. Purvis, and M. Purvis. Role model based mechanism for norm emergence in artificial agent societies. *LNCS*, 4870:203–217, 2008.

[5] O. Sen and S. Sen. Effects of social network topology and options on norm emergence. In *Proceedings of the 5th COIN*, COIN'09, pages 211–222, Berlin, Heidelberg, 2010. Springer-Verlag.

[6] Y. Shoham and M. Tennenholtz. On social laws for artificial agent societies: off-line design. *Journal of Artificial Intelligence*, 73(1-2):231–252, February 1995.