# Friends No More: Norm Enforcement in Multi-Agent Systems

Adrian Perreau de Pinninck Carles Sierra Marco Schorlemmer
IIIA - Artificial Intelligence Research Institute
CSIC - Spanish National Research Council
Bellaterra (Barcelona), Spain
{adrianp,sierra,marco}@iiia.csic.es

## ABSTRACT

We propose a new distributed mechanism to enforce norms by ostracizing agents that do not abide by them. Our simulations have shown that, although complete ostracism is not always possible, the mechanism substantially reduces the number of norm violations.

## Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent systems

## General Terms

Experimentation

## Keywords

Norm enforcement, Ostracism, Trust

## 1. INTRODUCTION

In a normative MAS a set of norms are defined to shape the set of available actions in order to accomplish the coordination between agents. An autonomous agent has the choice whether or not to support a norm. For a utility maximizer agent if following a norm is profitable, it is in the agent's own interest to act as the norm establishes. But this is not always the case, as some norms are only profitable when all agents abide by them. For this kind of norms, an agent that does not adhere (i.e., a violator) will profit at the expense of the agents that adhere.

The aim of this paper is to introduce a new distributed mechanism that attains norm compliance by ostracizing norm violating agents in an open MAS. Our scenario allows agents to interact with each other. An agent can interact with the agents it is linked to directly or indirectly through a path of links (i.e., agents can interact with direct neighbors, with neighbors of neighbors, and with their neighbors and so on...).


*AAMAS'07* May 14–18 2007, Honolulu, Hawai'i, USA.


We define interactions as a two-player game with two possible strategies; cooperate and defect. The utility function will be that of a prisoner's dilemma (see Figure 1).

| PD | Cooperate | Defect |
|---|---|---|
| Cooperate | 3,3 | 0,5 |
| Defect | 5,0 | 1,1 |

**Figure 1: Prisoner's Dilemma Payoff Matrix**

The norm in this scenario is for agents to cooperate with each other, thus attaining the maximum utility for the society. Nonetheless, agents can choose to violate the norm and defect. Violators are better off when interacting with a cooperative agent since they gain more utility. In order to attain norm enforcement, some agents (we will call them *enforcer* agents) are given the ability to stop interacting with violators, and to stop violators from interacting with their own neighbors. When enough agents use this ability against a violator, it becomes ostracized.

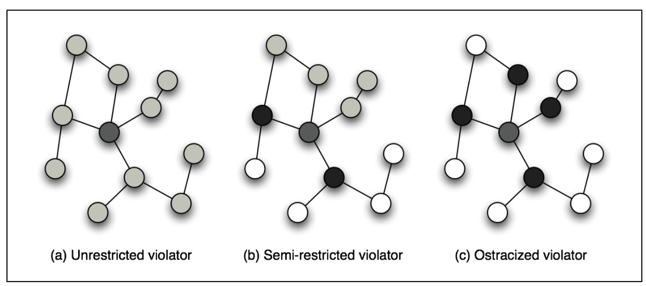


**Figure 2: Ostracizing a violator**

The ostracism process is illustrated in Figure 2. At first an undetected violator in the network (the dark gray node) can interact with all the other agents (light gray nodes are liable to interact with the violator). When the violator interacts, and defects, it can be detected by enforcer agents which will start blocking its interactions (black nodes are blocking agents, and white nodes are agents that the violator cannot interact with). When all the violator's neighbors block it, it is ostracized.

In Section 2 we describe the scenario we employ in the simulations. In Section 3 we analyze the simulation data.

## 2. THE SCENARIO

We model our multi-agent system as an undirected, irreflexive graph: $MAS = \langle Ag, Rel \rangle$, where $Ag$ is the set of vertices and $Rel$ the set of edges. Each vertex models an agent and each edge between two vertices denotes that the agents are linked to each other. In order for two agents to interact, there must be a path in the graph between the two.

An initiator agent will search for a path in the society to find a partner agent with which to interact. All the agents in the path that are not the initiator or the partner agent will be called mediator agents (i.e., agents mediating the interaction). Once an agent chooses a partner, a prisoner's dilemma game is played. The game results and the path are known by both playing agents. Playing agents can choose to send the game results to all the mediators in the path.

Our society of agents will be composed of three types of agents. *Meek agents* are norm-abiding agents that always cooperate. *Violator agents* which always defect. Finally, *enforcer agents* are norm-abiding agents that have the ability to block violators, which is essential in order to achieve their ostracism.

Two different blocking strategies have been tested. The Uni-Directional Blockage (UDB) strategy, where only known violators are blocked when searching for a partner. Or the Bi-Directional Blockage (BDB) strategy, where not only violators are blocked, but also norm-abiders are prevented from reaching known violators.

A simulation is made up of 100 agents arranged in one of three possible graph structures: tree, random, or small-world. Two parameters ranging from 0% to 100% in 10% increments define the quantity of enforcer and violator agents. Another parameter defines whether enforcer agents use the BDB or UDB enforcement strategy. Simulations are run during 1000 rounds. In a round each agent tries to find a partner, interacts through the prisoner's dilemma, and can gossip the results to the mediator agents. The average utility and norm violations received by each type of agent has been calculated in each simulation.

## 3. ANALYSIS

Figure 3 shows that the higher the percentage of norm-abiding agents that use a blocking rule, the lower the average number of norm violations received by any agent in our system. Figure 4 shows a larger reduction when looking only at the results for norm-abiding agents. There are eight different lines drawn in both graphs, each one stands for a different percentage of violating agents. In all cases a higher enforcer to meek agent ratio ($x$-axes) leads to lower violations received in average by any agent ($y$-axes).

We also deduce from the data that different organizational structures in the multi-agent system influence norm enforcement. In Figures 5 and 6 we have extracted the average norm violations ($y$-axes) for each of the different structures tested: Random, Small World, and Tree. The $x$-axes contains the different percentages of enforcer agents tested. It can be seen that both random and small world networks have an almost identical graph line. On the other hand the tree structure has shown to improve the enforcement capabilities.

In random and small world networks, when the percentage of enforcer agents reaches its maximum the percentage of violations received are increased (see Figure 6). We believe this happens because violator agents manage to form a sub-society, and in an interaction between two violator agents, two violations are being accounted for. This has been observed in simulations with a ratio of violator agents of 20% and above. When the ratio of violator agents is low enough, enforcers manage to block them completely.

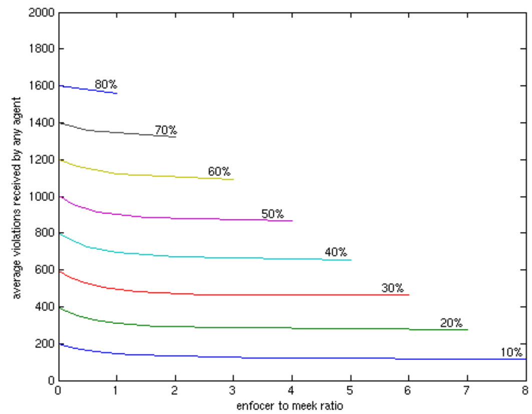


**Figure 3: Blocking reduces violations**

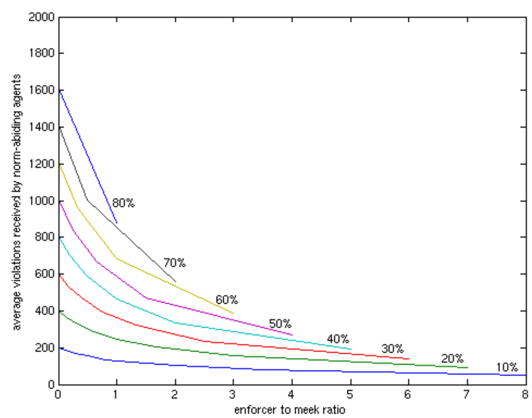


**Figure 4: Blocking reduces violations**

The data in Figure 7 shows that the enforcement strategy used by enforcer agents can reduce the number of violations received by meek agents. The $y$-axes contains the increment in efficiency at protecting meek agents from violations:

$$\Delta \mathit{Efficiency} = ((\mathit{Violations}_{BDB} / \mathit{Violations}_{UDB}) - 1) \times 100$$

We observe that for random and small world networks the efficiency is positively correlated with the enforcer to meek agent ratio. We can conclude that BDB has a higher efficiency at protecting meek agents from violator agents, unless violators are majority. This is not observed in the tree network, where the enforcement strategy does not have a significant influence. The tree network is already good enough for ostracizing offenders, and the BDB strategy does not improve that.

We assume a strategy is rational if it maximizes the agent's utility. What we have tested is whether following the norm maximizes the agent's utility. Figure 8 shows the utility gained ($y$-axes) by norm supporting agents, its $x$-axes shows

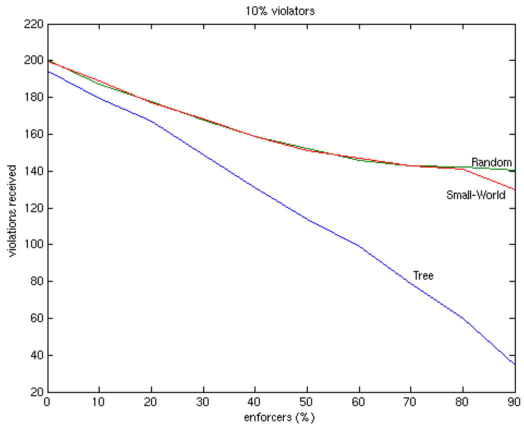


**Figure 5: Enforcement capabilities vary depending on structure (10% violators)**

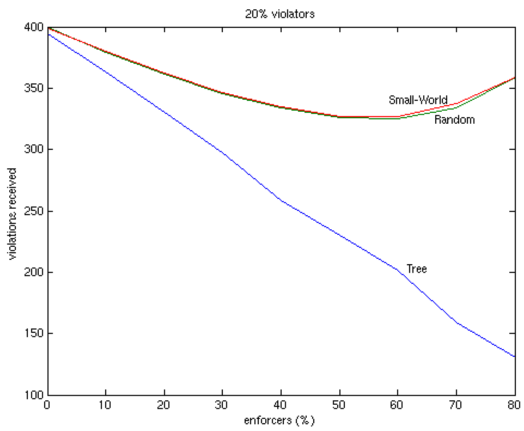


**Figure 6: Enforcement capabilities vary depending on structure (20% violators)**

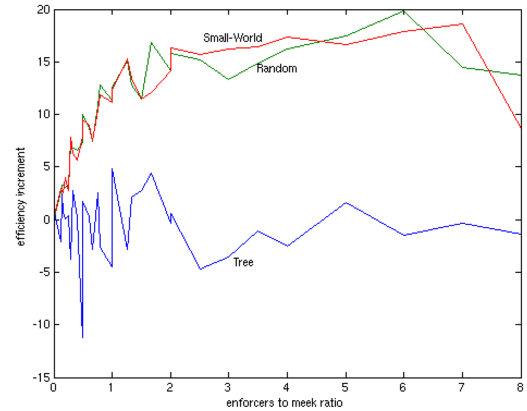


**Figure 7: Enforcement strategy influences received violations**

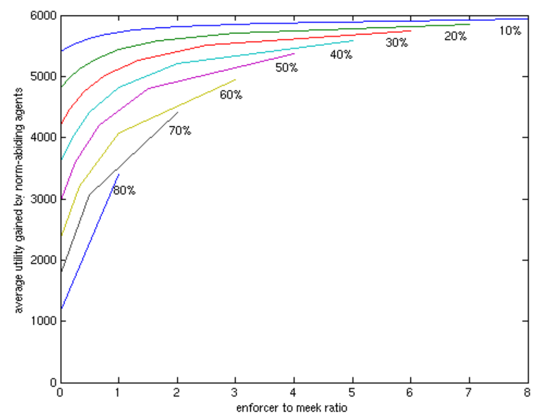


**Figure 8: Utility gained by norm-abiding agents**

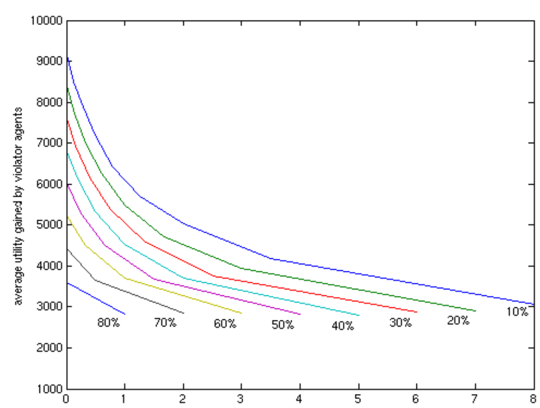


**Figure 9: Utility gained by norm-violating agents**

the enforcer to meek agent ratio. Each line stands for a different percentage of violating agents. Figure 9 instead shows the utility gained by norm violating agents. As the number of enforcers increases there is a tendency for norm supporters to gain more utility, while the opposite tendency is observed for violator agents. When the number of enforcer agents is low, the utility gained by violator agents is much higher than the one gained by norm supporters. As the number of enforcer agents grows the roles are reversed. The inflection point depends on the amount of violator agents in the system. Therefore we say that enforcement makes supporting the norm a rational strategy


## Acknowledgements

This work is supported by the OpenKnowledge[1] STREP project. A.Perreau de Pinninck is supported by an I3P - CSIC predoctoral fellowship. M.Schorlemmer is supported by a *Ramón y Cajal* research fellowship.


[1] `http://www.openk.org`