Citation for published version:

Balke, T, De Vos, M & Padget, J 2013, An agent-based simulation approach to comparative analysis of enforcement mechanisms. in *Coordination, Organizations, Institutions, and Norms in Agent Systems VIII: 14th International Workshop, COIN 2012, Held Co-located with AAMAS 2012, Valencia, Spain, June 5, 2012, Revised Selected Paper*. Lecture Notes in Computer Science, vol. 7756, Springer, Berlin, Germany, pp. 53-70, 14th International Workshop on Coordination, Organizations, Institutions, and Norms in Agent Systems VIII (COIN 2012), Valencia, Spain, 5/06/12. https://doi.org/10.1007/978-3-642-37756-3_4

DOI:

[10.1007/978-3-642-37756-3_4](https://doi.org/10.1007/978-3-642-37756-3_4)

Publication date:

2013

Document Version

Peer reviewed version

[Link to publication](#)

The final publication is available at http://dx.doi.org/10.1007/978-3-642-37756-3_4

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

An Agent-based Simulation Approach to Comparative Analysis of Enforcement Mechanisms

Tina Balke^{1,2}, Marina De Vos², and Julian Padget²

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. Incentive-based enforcement can be an effective mechanism for fostering cooperation in open distributed systems. The strength of such systems is the absence of a central controlling instance, but at the same time, they do depend upon (voluntary) regulation to achieve system goals, creating a potential “tragedy of the commons”. Many different mechanisms have been proposed, both in the multi-agent systems and the social science communities, to solve the commons problem by using incentive-based enforcement. This paper advocates the use of agent-based simulation to carry out detailed comparative analysis of competing enforcement mechanisms, by providing common settings, the environment and the basis for comprehensive statistical analysis. To advance this argument, we take the case study of wireless mobile grids, a future generation mobile phone concept, to ground our experiments and analyse three different enforcement approaches: police entities, image information and a well-known existing reputation mechanism. The contribution of this paper is not the enforcement mechanisms themselves, but their comparison in a common setting through which we demonstrate by simulation and statistical analysis that enforcement can improve cooperation and that a relatively small percentage (of the population as a whole) of police agents outperforms (under the chosen metrics) image- and reputation-based approaches. Hence, qualified conclusions may be drawn for the application of such mechanisms generally in open distributed systems.

1 Introduction

Open distributed systems allow autonomous entities with some form of social relationship to join and leave freely as well as to perform actions such as interacting with other entities. Entities base their decisions and actions on their own goals as well as their expectations about the system and the behaviour of the other entities. The result of the combined individual decisions and actions is a global emergent behaviour that—in contrast to the individual decision making processes—can be perceived from outside the system.

The principal advantage and disadvantage of open distributed systems is that at design-time, it is unknown precisely what individual and collective behaviour may be exhibited by participating entities. Complete control of even closed distributed systems has proven a very challenging problem. Rigid control of open systems, especially given

their increasingly pervasive nature, is unrealistic; not only is imposition of controls a reaction to a perceived threat (to system integrity), it also fails to recognize open systems as a nascent opportunity.

One particular problem in open distributed systems (be it relay-routing, peer-to-peer, cloud computing, etc.) is that they require some form of contribution on the part of participants, which translates into some form of cost to them. Participants can exhibit strategic behaviour and are not necessarily cooperating (i.e. contributing to the system). For an agent, making resources available therefore has the danger that its good behaviour is not reciprocated, resulting in no inherent value in cooperation for a participant. A lone cooperating user draws no benefit from their cooperation, even if the rest of the system might. Guaranteed cost paired with uncertainty or even lack of any resulting benefit does not induce cooperation for a utility-maximizing user. Without any further incentives, rational users therefore would not cooperate in such an environment and all will be worse off than if they cooperated. This phenomenon is referred to as the “Tragedy of the Commons” [14, 19]. As a consequence of the above problem, and of the limitations on the extent to which rigid control is feasible in open distributed systems, enforcement mechanisms offer a means to reduce the prevalence of the commons phenomenon.

To evaluate enforcement mechanisms empirically, as we propose here, in order to be able to identify which aspects can be quantified and how, a suitable domain is needed. However this inevitably creates the risk that decisions are made, or metrics are constructed that are domain-specific. The domain chosen in this paper is the wireless mobile grid (WMG), which is described in more detail in the next section. We do not make a judgement as to whether the WMG concept is viable or not: it is simply a novel example of the kind of emerging ‘digital commons’ that makes it suitable as a case study for which it would also be useful to get some early indicators of which kinds of enforcement are effective and what the associated costs might be.

The WMG, as an opportunistic network made possible by chance co-location, exhibits many of the characteristics of an open system: participants are free to join or leave at any time, identity is not authenticated and free-riding appears to be easy. Consequently, repeat encounters are likely to be few and participant turnover high. Thus, the participant contributions required to sustain it may be difficult to acquire or to incentivize. The relative complexity of the scenario, at least until better understood, makes an analytical or game-theoretic approach infeasible at this stage, so we advocate agent-based simulation as a means to establish a better understanding of the dynamics and to evaluate side-by-side three well-established enforcement mechanisms. Based on the common setting provided by the simulation environment we examine the mechanisms’ advantages and disadvantages in respect of one another. For this purpose, in the next section, we describe the case study, then in Section 3 we present the three enforcement mechanisms to be compared. The simulation experiments and its results are discussed in Sections 5 and 6. This paper closes with a short summary of the findings as well as a discussion of their implications for open distributed systems (Section 7).

2 The Wireless Mobile Grid Case Study

To demonstrate the use of agent-based simulation for the comparative analysis of enforcement mechanisms, we start by establishing a common case study for our experiments which portrays the particular features of open distributed systems sketched in the previous section. In one sense, the domain details of the case study are not especially important, but are simply there to ground the scenario, rather than using an abstract scenario which can be harder to assimilate. Thus, the particular domain is the so-called “wireless mobile grid” (WMG); a mechanism proposed by Fitzek and Katz [10] to address the energy issues inherent in 4th generation mobile phones. This paper is not about the plausibility or otherwise of WMGs, but about the comparative effectiveness of different approaches to enforcement in the context generated by WMGs, representing an instance of the broader class of open distributed systems.

In WMGs, as well as using the traditional 3G (or LTE) communication link with base stations, users are envisioned as sharing resources in a peer-to-peer fashion using a short-link connection protocol such as IEEE802.11 WLAN. The advantage of this short-link connection is that it uses less power and allows for higher data rates. However, in order to function properly WMGs require collaboration between users, which may be difficult to realize. The main problem in WMG is that collaboration comes at the cost of further battery consumption. In consequence, rational users will prefer to receive the resources without any commitment to contribute themselves. However, if a substantial number of users follow this selfish strategy, the WMG will not work and none will benefit from the potential energy savings arising from cooperation [25]. A WMG is an interesting and novel example of an open distributed system, in which there are autonomous users with their own goals who can freely join and leave, who can interact with one another over short-range connections for short periods, and whose individual actions contribute to the success or failure of the WMG as a whole.

Purely technical (hard-ware or hard-coded) solutions for ensuring non-compliance in open distributed systems are frequently subverted (see [17] for example). Hence, the approach we take here, which is to employ enforcement mechanisms such as reputation information or police agents that regulate WMGs by social means.

3 Enforcement Mechanisms for Wireless Mobile Grids

Many possible enforcement mechanisms exist, so how to choose suitable ones? Balke and Villatoro [3] provide a systematic overview of possible enforcement options by identifying the roles actors can have in an enforcement setting and discussing all possible combinations of these roles in the enforcement process. We do not reproduce the details of the mechanisms analysed in [3], for sake of space as much as correctness, but draw on their conclusions to select three mechanisms to concentrate on in more detail: reputation information, image information, also known as direct trust, and police agents. We choose these because of their popularity in the agent community as much as for their complementary foundational concepts that allows us to explore a wider range of options.

3.1 Utilization of Police Entities

The utilization of police entities can be thought of as the implementation of entities with normative power (e.g. some kind of policing) [15] that participate in the system (in our example, the WMG) and have permission from the system's owner to punish, if detected, negative/inappropriate behaviour (i.e. non-compliance) by means of sanctions. In contrast to regimentation (i.e. complete control) [7], the police entities do not control all actions but only act as enforcers when violations are detected. Detection of violations is done by the police entities themselves, who test the behaviour of entities and react to what they detect. Several kinds of sanctions can be imagined depending on the severity of the non-compliance, such as complete exclusion from the WMG or penalty payments either monetary or in terms of energy.

3.2 Image Information

Image information [18], also called direct trust, is a global or averaged evaluation of a given agent – usually called the target – on the part of an individual. It consists of a set of evaluative beliefs about the characteristics of a target. These evaluative beliefs concern the ability or the possibility of the target to fulfil one or more of the evaluator's goals, e.g. to cooperate in a WMG transaction. An image basically gives the evaluator's opinion of whether the target is “good” or “bad” or “not so bad” etc. with respect to a norm, a standard, a skill etc. When utilizing image information, an agent uses *its own* information about the past behaviour of the potential interaction partner and makes decisions based on this information.

3.3 Reputation Information

Reputation information, in contrast to image information, comprises not only agent's own acquired image information, but that obtained from other agents as well. Thus, reputation in this paper is understood as the process of and the effect of the transmission of a target's image. In contrast to image information alone, as described above, when images are circulated more information becomes available to the individual agents. However, the circulation of information itself can generate costs. Furthermore it is possible that agents may circulate false image information to increase their value relative to other agents.

4 Related Work

Looking at previous works that are of importance for this paper, one can look into two different directions. The first direction is related work on means of enforcement in open distributed systems, such as WMGs, whereas the second direction is related work dealing with the comparison of enforcement mechanisms.

Looking at the first direction, one can identify a large literature in economics and social sciences on cooperation and free-riding and the mechanisms to overcome the latter. One of the most well-known analyses is by Ostrom [20], who shows that in small

(relatively closed) communities these “tragedies” can be overcome. Other works use game theory [4] or evolutionary game theory [12, 13] to address the question. In general, most of the works looking into enforcement use some form of punishment which serves as a deterrent to the rational behaviour of utility-based participants (e.g. [9]). Thus, a punishment is a fine taken from the participant’s benefits. The topic has been framed mostly in terms of mechanism design and the issues that economists have studied more thoroughly are the information about infraction and sanctions [8], as well as the amount and pervasiveness of sanctions [6]. As pointed out before methodology often is either (evolutionary) game-theoretic (see [4] for example) or experimental (including agent-based simulations) [11, 22].

In the second direction, i.e. the comparative study of enforcement, little related work can be found. In [5] for example, the authors discuss differences between image and reputation in detail, however no detailed experiments testing their impact on the same setting are made. Similar in [22] the authors test the impact of how far messages are sent in a network and even considers the costs of these messages, but no comparison between trust and reputation is made.

5 The Simulation Design

Having briefly outlined the enforcement mechanisms under examination, we now present the basic simulation setup. We first describe the agents and their decision making behaviour and then outline how the enforcement mechanisms have been implemented. Concerning the technical components of the wireless mobile grid, we adopt the well-established “flat earth model” [16] that assumes symmetry (i.e. if node A can hear node B , B can hear A) and an absence of obstacles that might reduce transmission quality. The flat-earth model is a widely accepted simplification made in the mobile communications community and has been used for simulation presented in mobile communication centred articles on this topic (e.g. [2]). Furthermore we assume that all agents have identical mobile phones, for which we use the energy consumption profile data reported in [21]. A reason for this assumption is that [21] explains that the Nokia N95 mobile phone is a representative phone with features for WMG communication and that the differences between different mobile phones are only marginal.

5.1 The Basic Agent Decision Process

The simulation uses one agent for each user/mobile phone pair. These agents move randomly in the simulation space and at any given point of time can interact with the agents that are within their (modelled) WLAN range. The agents make decisions that maximize their utility under the constraint of bounded rationality. Different agents are given different utility valuations. We define three kinds of non-police agents according to the behaviours for which they maximise: (i) “utility agents” that try to minimize battery consumption and avoid punishment (ii) “honest agents” that cooperate whenever possible, and (iii) “malicious agents” that try to undermine the system regardless of cost. The agent’s decision-making is based on incomplete knowledge of the system state, so they can only optimize for local utility, which may be different from the global utility.

Local knowledge is determined by two factors: the agent's location and its WLAN radius. A full Cartesian model is unnecessary, since we only need to model proximity, hence an agent location is modelled as $l \in \mathbb{R} \bmod 1$, that is the interval wraps around. An agent at ϵ and another at $1-\epsilon$ are 2ϵ apart. The proximity of two agents is determined by each agent's WLAN radius (r_v). An agent at l_1 has radius $[l_1 - r_{v1}, l_1 + r_{v1}]$ and another at l_2 has radius $[l_2 - r_{v2}, l_2 + r_{v2}]$. Communication between these two is possible if these intervals intersect.

The procedure for the agent's decision making process and its utility considerations (see Figure 1) are based on the following issues:

1. Each agent has the task of acquiring a whole file through downloading (over 3G) or exchanging (over WLAN) file chunks. The agent must decide whether to download it all or to search for a collaboration partner with whom to share the work. File size is the first determinant: if the file is small and the potential costs of finding a collaborator are higher than the potential gains, then the agent will download all the chunks itself. Otherwise, it looks for nearby agents.
2. If the neighbourhood is sparsely populated, then the chances of finding sufficient partners is low and the agent downloads all the chunks itself. Otherwise, it sends a cooperation request specifying the file whose chunks it requires using WLAN broadcast.
3. If the agent receives a cooperation request for the same file, there is no need to send a request, so it just replies using WLAN broadcast.
4. Having sent a cooperation request, the agent awaits responses. From the positive replies, the agent selects collaboration partners, possibly using image or reputation mechanisms to decide.
5. Once a cooperation group has been formed, an agent has to decide: (i) whether to download its promised chunk(s) (over 3G) and (ii) whether to share its chunks with the cooperation group.
6. The cooperation decision depends on the agent's individual utilities for cooperation and defection. Thus, an utility agent (see definitions above) compares energy benefits from defecting now, against the future costs arising from detection in terms of the likelihood and level of a fine, by comparing the number of past defections with the number and level of past fines. An honest agent will always cooperate and will never defect. A malicious agent, in contrast, will always defect.

Having made its decision, the last step is to wait and see whether the cooperation partners send their promised shares. For missing shares, the agent repeats the decision process outlined above.

5.2 Implementing the Enforcement Mechanisms

We now describe the implementation of the three enforcement mechanisms.

Enforcement Agents An police agent has the same properties as a ordinary agents, except for restricted behaviour in that it: (i) responds positively to cooperation requests, if not already committed, and then performs its share of downloading and sharing,

Fig. 1. Download considerations

(ii) does not have its own file download tasks set at the beginning and (iii) does not send cooperation requests.

A police agent never defects. Its energy costs count towards the total WMG energy costs – so that enforcement costs are included. It also monitors cooperation by checking on the number of defections in the cooperation group to which it belongs and fines those that have defected. The fine is measured in terms of energy and is set at three times the agent's relative gain from defection³.

Enforcers make no distinction between intention and absence of action: it only matters whether an agent kept the cooperation agreement by sending its chunks by the cooperation deadline. The same defection may be observed by more than one enforcer, but the fining mechanism ensures that an agent is not fined twice for the same offence.

Image Information Image information is used by agents in making cooperation decisions. No police agents are used. Image information is information acquired about another agent through direct interaction with that agent. Thus, the experience of past interactions is used to evaluate new cooperation requests. If the image is positive, the collaboration proceeds. Each agent could have an individual defection tolerance level for all its past collaborators, but to keep the range of experiments within what can be reported in this paper, this is uniformly set to zero. Thus, the object of a defection will never collaborate with the defector again. If an agent has no image for the originator of a cooperation request, it treats the image as positive.

The next step is to incorporate image information in the utility computation. First, the agent computes the fraction of the number of times its cooperation (request and offer) was rejected because of its image. This value is then weighted by (i) the chunk size and (ii) the cooperation group size. The latter is significant because, the larger the group, the greater the spread of negative image information if the agents defect.

The advantage of image information is its reliability. However, one problem often associated with image mechanisms is that an agent first needs its own experience to construct image information. Consequently, it can always be the object of a defection at least once. Reputation mechanisms are proposed as a way to avoid this problem. We next explore this approach.

Reputation Information In reputation mechanisms, the image information of individual agents is circulated. A large number of reputation mechanisms have been put forward in the literature covering a range of circumstances. To select one suitable for WMG, we start by examining the requirements and constraints.

Although the work reported here is simulation-based, in a real WMG, the actual agents will include humans. This suggests any mechanism should allow for the subjective expression of trust based on individuals' perceptions. Additionally, the mechanism needs to be compatible with the non-numerical and non-monotonic models of human expression. Finally, the mechanism must be able to handle incorrect information, taking

³ We tested a range of alternative fine levels. Space limits prevent a full report, but the value of three times the gain exhibited the best balance between deterrent and the fine not being disproportionate to the offence.

the sources of (reputation) information into account and identifying those that provide false information.

Consequently, we consider three candidate mechanisms: Regret, Fire and Abdul-Rahman and Hailes (ARH), which are analysed in detail in [23]. Regret seems unsuitable because by default it requires a large number of messages to be sent (accounting for witness, neighbourhood and system reputation) which inevitably increases overall energy consumption. Fire is also unsuitable because of its basic assumptions that agents (i) willingly share their experience and (ii) report truthfully when exchanging information with one another. This leaves the ARH reputation mechanism [1], which fortunately satisfies our requirements.

ARH requires that each agent maintains a database of trust relationships that they use for themselves or to respond to the requests of others. The data is segregated into direct (image) and indirect (reputation) information. ARH defines a “trust-relationship” as a vectored connection between exactly two entities, which in some circumstances can be transitive. In this way they distinguish between direct trust relationships (“Alice trusts Bob.”) and recommender trust relationships (“Alice trusts Bob’s recommendations about the trustworthiness of other agents”). This allows entities to account for the source of reputation information as well as collecting and evaluating information about the reliability of recommenders. Another interesting contrast to other formalizations is that, reflecting the qualitative nature of trust, ARH does not use probability values or the $[-1, 1]$ interval, but a multi-context recording model with abstract trust categories that are easier for humans to understand. These trust values relate to certain contextual information (“Alice trusts Bob, concerning “table”-transactions. However, she does not trust him when it comes to “chair”-transactions.”).

ARH models trust as context-dependent, so it is defined as a “troika” of (*agent-ID*, *Trust-Category*, *Trust-Value*), with trust categories such as “cooperation partner” or “recommender”. ARH sets out a recommendation protocol for handling recommendation requests, statements and enquiries. A recommendation request is forwarded until one or more agents are found that can give information for the requested category and which is trusted by the penultimate agent in the chain. We do not support routing in the WMG simulation and implementing the protocol described above would result in large amounts of network traffic, impacting significantly upon the potential benefits of WMG. Thus, we simplify this aspect of ARH in such a way that an agent seeking a recommendation about a target sends out *one* broadcast message to its neighbours. If it receives no answers, the agent does not wait for further information, but as in the case of image information, cooperates with the target.

An agent, of whichever kind, uses the reputation mechanisms as follows:

1. If it has image or reputation information about the potential cooperation partner, it uses it.
2. If not, it sends a request for recommendations to its neighbours:
 - (a) If there are no replies, the agent agrees to collaborate and will update its image information in due course in respect of the outcome of the collaboration.
 - (b) If there are one or more replies, they are categorized by source into trusted, untrusted and unknown:

- i. Trusted source: the agent updates its local reputation information using the most trusted source and acts accordingly (i.e. positive recommendation: collaborate, negative: not). For equally trusted sources, the first reply is used.
- ii. Untrusted source: the information is kept for later validation but not taken into account for the current decision.
- iii. Unknown source: information is treated as for a trusted source.

Different kinds of agent respond differently to reputation requests. The utility maximizing agent will not send any information, because answering a message costs energy. An honest or a malicious agent answers on average one request per interaction event, in order to limit energy spent on answering reputation requests. An honest agent always reports truthfully about the target, with the aim of improving the overall information level in the system. A malicious agent however, if the target is not itself, always gives negative feedback on the target, with the aim of enhancing its relative reputation.

6 Simulation Setup and Results

6.1 Simulation Setup

As pointed out in the introduction this paper focuses on the use of agent-based simulation to carry out a detailed comparative analysis of competing mechanisms (i.e. enforcement mechanisms in our scenario) and to determine which of these meets the systems objective (which we defined as energy saving) best. To test the impact of the three different enforcement mechanisms on the cooperation problem and the resulting energy consumption in WMGs, we formulate the following hypothesis:

Hypothesis 1: The presence of an enforcement mechanisms reduces the average energy consumption compared to when there is none.

We can assume that an enforcement mechanism has an effect on energy consumption, but if it affects cooperation – and in consequence energy consumption – is this effect constant across various simulation settings? Specifically, we want to establish sensitivity across a range of parameters: (i) population size, $|\mathcal{A}|$ (ii) population density (the average number of agents within each others' WLAN radius), $\rho_{neighbourhood}$, and (iii) population composition, that is proportions of utility, honest and malicious agents.

To test the influence of $|\mathcal{A}|$ and whether either of the other two parameters affects the simulation results, we check the null-hypotheses that no difference in simulation results can be observed when these parameters are varied. We then examine what impacts upon the different enforcement mechanisms:

Hypothesis 2: The success (in terms of the average energy consumption) of a WMG using reputation-based enforcement depends on population size, density and composition.

Hypothesis 3: The success (measured by average energy consumption) of a WMG using police agents as the enforcement mechanism depends on population size, density and composition as well as the number of police agents $|\mathcal{A}_{\text{Enf}}|$.

Table 1. Simulation Variables

Name	Range/Type	Simulation Parameter
Number of Agents ($ \mathcal{A} $)	$[2, \infty]$	200, 400, 800
Utility Agents as % of $ \mathcal{A} $	$[0,100]$	0, 25, 50, 75, 100
Malicious Agents as % of $ \mathcal{A} $	$[0,100]$	0, 25, 50, 75
Honest Agents as % of $ \mathcal{A} $	$[0,100]$	0, 25, 50, 75
Enforcement Mechanism		None, Police Agents, Image Info., Reputation Info.
Number of Police Agents $ \mathcal{A}_{\text{Enf}} $ as % of $ \mathcal{A} $	$[0,\infty]$	0.5, 1, 2, 3, 5 % of $ \mathcal{A} $, s.t. $ \mathcal{A}_{\text{Enf}} > 1$
$\rho_{\text{neighbourhood}}$	$[0, \mathcal{A} - 1]$	10, 20

Table 2. Analysis of variance of experiments with and without enforcement

Source	Sum of Squares	Degrees of Freedom	Mean Squares	F	Prob > F (= p-value)
Enforcement	224.759	4	56.1898	1539.94	< 0.0001
Error	709.842	19454	0.0365		
Total	934.601	19458			

Hypothesis 4: The success (measured by average energy consumption) of a WMG using image-based enforcement depends on population size, density and composition.

For all of the above, we use the experiment configuration shown in Table 1, which summarizes the factorial experiments performed and the values over which each simulation parameter ranges.

6.2 Simulation Results

The experiments consist of 50 runs for each of the 468 parameter combinations in Table 1, making 23,400 runs in total. We used ANOVA to test the significance relationship between the independent variables (the parameters in the simulation) and the dependant variables (the number and ratio of defections and energy consumption)⁴. We also applied Tukey’s test as a post-hoc ANOVA, which identifies the impact of specific variables on the overall result.

Testing Hypothesis 1 We can now analyse the simulation results to test the hypotheses formulated in the previous section. First, we test hypothesis 1 and look at mean energy consumption when there are different enforcement mechanisms employed. By means of ANOVA, we can test whether there is sufficient evidence to reject the null hypothesis that enforcement mechanisms have no effect on energy consumption. Table 2 shows the results of this comparison.

⁴ We performed the Shapiro-Wilk test and Levene’s test to ensure the applicability of ANOVA. Due to limits on space we do not include details of these results

Fig. 2. Multiple Comparison (Tukey's Test) Results of Marginal Means for Comparing Simulation Experiments with and without Enforcement – Post Hoc Test

Table 3. Post-hoc analysis of variance of enforcement mechanisms

Source	Sum of Squares	Degrees of Freedom	Mean Squares	F	Prob > F (= p-value)
Image-Information	26.969	1	26.969	684.73	< 0.0001
Police Agents	168.95	5	33.7907	801.85	< 0.0001
Reputation	12.308	1	12.3078	332.01	< 0.0001

As the significant p-value suggests, we can reject the null hypothesis and conclude that the utilization of enforcement results in a difference in the average energy consumption. Looking at the parameters that influenced this result the most, we see that the success of enforcement mechanisms was significantly dependant on the population composition ($p < 0.0001$) which makes our Hypotheses 2–4 correct with respect to that parameter. Analysing Table 2 more closely, one notices that a high error rate can be observed, indicating that a difference exists between the three enforcement mechanisms that are grouped in the ANOVA. In order to examine this effect more closely, as well as to determine the extent to which each enforcement mechanism contributes to this difference, we perform Tukey's test as post-hoc analysis. Fig 2 shows the results of this analysis and Table 3 gives an overview of the respective statistical values per enforcement mechanism.

As the p-values in Table 3 show, all mechanisms have an average energy consumption significantly different to the experiments with no enforcement, however looking at the Tukey's test results in Fig. 2, it is clear that the results for the reputation mechanism stand out. Thus, whereas the results indicate that we can confirm hypothesis 1 for image-related information and police agents, the Tukey's test for the reputation mecha-

Fig. 3. Multiple Comparison (Tukey's Test) Results of Population Size Marginal Means in Settings with Reputation Information – Post Hoc Test

nism show that the experiments using this mechanism, have an on average higher mean than when no enforcement is used. This implies that the utilization of reputation information *increased* the average energy consumption.

Testing Hypothesis 2 Digging deeper into the differences in energy consumption, the reason for this effect becomes apparent. As a result of the large number of additional messages arising from the transmission of reputation information, the communication costs are disproportionately high and thereby increase energy consumption. Thus, especially in cases with large numbers of honest agents, which would choose cooperation even without enforcement mechanisms, the reputation requests and answers do not improve enforcement, but rather result in additional energy consumption. This leads to reputation being worse with respect to the overall energy consumption ratio even compared to image information, i.e. settings where agents could only rely on their own personal experiences.

This effect from the additional communication costs can also be seen when looking at the ANOVA and Tukey's test result for $\rho_{neighbourhood}$ (Figure 4), which show that a higher $\rho_{neighbourhood}$ tends to result in worse energy consumption⁵. This can again be attributed to the increased number of cooperation messages in these settings (with a higher $\rho_{neighbourhood}$ more agents receive and send messages). These additional message costs even outweigh the benefits of being able to observe more agents (because of the higher number of neighbours in the system). A second effect that impacted the reputation mechanism was the negative information introduced by malicious agents. Our simulation experiments were set up in such a way that if in doubt (i.e not verifiable), reputation information was considered to be correct (i.e. information from unknown sources was considered correct at the beginning). As a result of this, especially at the start of each experiment, by giving negative reputation information, malicious agents

⁵ The p-value for this relation is 0.3052, i.e. it is not significant. Nevertheless a tendency towards the described effect was detectable throughout the experiments.

Fig. 4. Multiple Comparison (Tukey's Test) Results of Neighbourhood Density Marginal Means in Settings with Reputation Information – Post Hoc Test

were able to discourage agents from cooperating with potential rivals, which also had the side-effect that agents could gather less image information of their own. As a consequence, they were more likely to have to trust uncertain reputation information in the following interactions, causing problems for the overall reputation mechanism. As a result of these findings, we conclude that the reputation mechanism chosen for our particular case study is unsuitable. The reasons for this however seem of a more general nature, that is, they are applicable to reputation mechanisms in general and should be considered when thinking about employing a reputation mechanisms for enforcement purposes: (i) Transmission of reputation information costs, for both sender and receiver. These communication costs might result in reduced contribution of information, and the costs associated with it might outweigh the benefits of the system. (ii) In our experiments false information clearly harmed the system. A reputation mechanism therefore needs to be able to cope with false information. (iii) Most reputation mechanisms rely on small communities in order to function well. For space limitations, we did not present the detailed figures, but our experiments indicate that an increase in the population size is deleterious for the reputation mechanisms, as more messages are sent and agents are less likely to encounter one another again soon.

Testing Hypothesis 3 A second interesting effect we can observe in Fig 2 is that an increase in the percentage of police agents does *not* result in a reduction of the average energy consumption ratio. Earlier on, in Table 3 we established that police agents can contribute to energy consumption reduction in WMGs, however exploring further, the value of this statement varies significantly with the percentage of police agents being employed (this is also indicated by a relatively high sum of squares error of 985.84 for the police agent value in Table 3). In the figure, the average mean energy consumption for experiments with 1%, 2% and 3% police agents are lower than the one with 5% and testing this result for significance, one finds significant evidence against the null hypothesis that the means of the results with 1%, 2% and 3% police agents are not

smaller than the results with 5% police agents (respective p-values < 0.0001). This implies that we have to reject the null hypothesis, which in turn means that the lower average energy consumption values for results with 1%, 2% and 3% police agents are not a result of chance. Comparing the number of defections in the experiments with police agents by performing a t-test, only a slight, and not significant, advantage for experiments with 5% police agents can be found, i.e. that in these settings police agents only added slightly to the total energy consumption. This implies that the improved detection of violations resulting from the larger number of police agents, is outweighed by the additional energy they consume. In economic terms this means that the lower percentage of police agents performs better with regard to satisficing cooperation when considering energy consumption. In economics, *satisficing* refers to a decision-making strategy that attempts to meet criteria for adequacy, rather than to identify an optimal solution [24]. Thus, although not optimal with regard to the detection of violations (1%, 2% and 3% police agents will detect less than 5%) the costs associated with them (i.e. the energy they consume for performing their observation and punishing actions) are significantly lower, making them more advantageous in terms of the overall energy saving.

Concerning the remaining parameters addressed in hypothesis 3, we found significant evidence that the null-hypotheses (i.e. that there is no impact on the parameters) can be rejected both for population composition and for neighbourhood density (respective p-values < 0.0001), while the significance levels for the size of the population varying between (0.0154 and 0.0604) this does not allow rejection of the null-hypothesis at a significance level of 0.01, but still indicates that there is reason to believe that population size could have a moderate impact.

Testing Hypothesis 4 For simulations in which image information is used, as hypothesized, we can reject the null hypothesis for all input parameters (i.e we have enough evidence to assume that Hypothesis 4 is correct). Both the population composition and the population size have a significance level $p < 0.0001$ and $\rho_{neighbourhood}$ has a p-value < 0.0005 .

Summarizing the Findings Summarizing for the four hypotheses formulated in Sec. 6.1 we arrive at the following conclusions:

Hypothesis 1: Correct for police agents (between 1% and 5% of $|\mathcal{A}|$) and for image information, but incorrect for reputation information.

Hypothesis 2: Reputation information did not help to decrease the average energy consumption in the WMG. Both population size and composition had a significant impact here.

Hypothesis 3: The success of police agents is dependent on population density, composition and the number of police agents $|\mathcal{A}_{\text{Enf}}|$.

Hypothesis 4: Correct for all parameters.

7 Conclusions

We have compared three different enforcement mechanisms that are often employed in open distributed systems in the context of a single case study, namely WMGs. WMGs broadly exhibit some of the most challenging characteristics of open systems, in the form of (potentially) rapidly changing participation and little or no authentication. Perhaps the most contentious aspect of the measurements, in respect of generalization, is the focus on energy consumption, since this is clearly specific to the WMG domain. The question is to what extent can this be viewed as a proxy for the fitness of an open system. We do not pretend that it is an accurate such indicator, but being an aggregator of interactions within the open system, and its minimization being a metric for the effectiveness of the enforcement mechanisms, it seems likely to be positively correlated with overall system fitness. A second issue affecting broader applicability may be the costs attributed to enforcement, since these are unavoidably expressed in domain-specific terms. Nevertheless, as with the previous point, the costs are just a formula associated with a transaction, and while changing the formula may lead to a different preference outcome, it should not fundamentally affect the validity of the approach.

Thus, the primary and unsurprising conclusion of this comparative analysis is that enforcement does not always help, and that the costs of enforcement need to be accounted for when deciding upon an enforcement mechanism, whether intrinsic in the form of, say communications, or extrinsic, in the form of, say rewards for observers looking for infractions. Of the three mechanisms presented, police agents especially seem to help to improve energy consumption. Although the results we present are inevitably linked to the specific mechanisms chosen, some general findings can be made.

The first is that the population composition accounts for the majority of the impact by the input factors on energy consumption. Second, the costs associated with the enforcement can outweigh the benefits. In the case of police agents this resulted in the situation that fewer agents produced a better absolute result in terms of energy consumption, while “only” satisficing the detection of violation actions. Similar effects could be seen in experiments with reputation information, where the message overhead produced by the reputation request and answers outweighed the benefits of the mechanism. One further aspect that influenced the performance of our reputation information-based enforcement mechanism negatively, is false information. This is particularly important as we implemented an adaptation of a mechanism that tried to account for this problem. However it did not have sufficient interactions to have any significant effect. The mechanism of Abdul-Rahman and Hailes – like any reputation mechanism – works better when the number of repeated interactions increases. Unfortunately this seems unlikely in a WMG and in open distributed systems in general, suggesting that the designer needs to consider – and possibly evaluate – carefully whether the characteristics of the situation are compatible with reputation mechanisms.

With respect to future work we aim to extend the simulation experiments by employing combinations of different enforcement mechanisms. In the experiments leading to this paper, we assumed that only one enforcement mechanism can be employed at any given time and that the choice of mechanism is constant throughout a simulation experiment. We plan to extend the work by relaxing this assumption and combining

different enforcement concepts in the simulation experiments. This paper established the baseline for each mechanism in isolation.

The experiments could also be extended by employing more sophisticated agents. We employed three kinds of agents that pursue very different strategies in order to test how sensitive the simulation is to very one-sided behaviour (e.g. always defect or always cooperate). In an actual deployment of a WMG such a one-sided behaviour might not be very realistic. Therefore, agents with more sophisticated reasoning processes that exhibit more diverse responses to the successes or failures of cooperation situations are needed. One extension could be to allow malicious agents to cooperate occasionally in order to make them harder to detect for other agents, or to allow for variations in the reactions to sanctions by the utility maximizing agents.

Acknowledgments

Tina Balke was partially supported by funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 288147.

References

1. Abdul-Rahman, A., Hailes, S.: Supporting trust in virtual communities. In: HICSS (2000)
2. Balke, T., De Vos, M., Padget, J.A.: Analysing energy-incentivized cooperation in next generation mobile networks using normative frameworks and an agent-based simulation. *Future Generation Computer Systems Journal* 27(8), 1092–1102 (October 2011), <http://www.sciencedirect.com/science/article/pii/S0167739X11000574>
3. Balke, T., Villatoro, D.: Operationalization of the sanctioning process in hedonic artificial societies. In: Workshop on Coordination, Organization, Institutions and Norms in Multiagent Systems @ AAMAS 2011, Taiwan (2011)
4. Coleman, J.S.: Foundations of social theory (August 1998), <http://www.amazon.ca/exec/obidos/redirect?tag=citeulike09-20&path=ASIN/0674312260>
5. Conte, R., Paolucci, M.: Reputation in Artificial Societies: Social Beliefs for Social Order. Springer (October 2002), <http://www.amazon.ca/exec/obidos/redirect?tag=citeulike09-20&path=ASIN/1402071868>
6. Dreber, A., Rand, D., Fudenberg, D., Nowak, M.: Winners don't punish. *Nature* 452, 348–351 (2008)
7. Esteva, M., Rodríguez-Aguilar, J.A., Sierra, C., Garcia, P., Arcos, J.L.: On the formal specifications of electronic institutions. In: Agent Mediated Electronic Commerce, The European AgentLink Perspective. pp. 126–147. Springer (2001)
8. Fehr, E., Gächter, S.: Cooperation and punishment in public goods experiments. *The American Economic Review* 90(4), 980–994 (2000), <http://dx.doi.org/10.2307/117319>
9. Feldman, M., Papadimitriou, C., Chuang, J., Stoica, I.: Free-riding and whitewashing in peer-to-peer systems. In: Proceedings of the ACM SIGCOMM workshop on Practice and theory of incentives in networked systems. ACM (2004)
10. Fitzek, F.H.P., Katz, M.D.: Cellular controlled peer to peer communications: Overview and potentials. In: Fitzek, F.H.P., Katz, M.D. (eds.) *Cognitive Wireless Networks*, pp. 31–59. Springer (2007)

11. Gurerk, O., Irlenbusch, B., Rockenbach, B.: The competitive advantage of sanctioning institutions. *Science* 312(5770), 108–111 (April 2006), <http://dx.doi.org/10.1126/science.1123633>
12. Güth, W., Ockenfels, A.: Evolutionary norm enforcement. *Journal of Institutional and Theoretical Economics* 156(2), 335–347 (2000), <http://edoc.hu-berlin.de/series/sfb-373-papers/1999-84/PDF/84.pdf>
13. Güth, W., Ockenfels, A.: The coevolution of trust and institutions in anonymous and non-anonymous communities. Discussion Papers on Strategic Interaction 2002-07, Max Planck Institute of Economics, Strategic Interaction Group (March 2002), <ftp://papers.mpiew-jena.mpg.de/esi/discussionpapers/2002-07.pdf>
14. Hardin, G.: The tragedy of the commons. *Science* 162, 1243–1248 (1968), http://www.garretthardinsociety.org/articles/art_tragedy_of_the_commons.html
15. Jones, A.J.I., Sergot, M.J.: A formal characterisation of institutionalised power. *Logic Journal of the IGPL* 4(3), 427–443 (1996), http://www-lp.doc.ic.ac.uk/_lp/Sergot/InstitPower.ps.gz
16. Kotz, D., Newport, C., Gray, R.S., Liu, J., Yuan, Y., Elliott, C.: Experimental evaluation of wireless simulation assumptions. In: Proceedings of the 7th ACM international symposium on Modeling, analysis and simulation of wireless and mobile systems. pp. 78–82. ACM, New York, NY, USA (2004), <http://users.cis.fiu.edu/~liux/research/papers/axiom-mswim04.pdf>
17. Leibowitz, N., Ripeanu, M., Wierzbicki, A.: Deconstructing the kaza network. In: Proceedings of the Third IEEE Workshop on Internet Applications. IEEE Computer Society (2003), <http://portal.acm.org/citation.cfm?id=832311.837393>
18. Miceli, M., Castelfranchi, C.: The role of evaluation in cognition and social interaction. In: Dautenhahn, K. (ed.) *Human cognition and social agent technology*. Benjamins, Amsterdam (2000)
19. Ostrom, E.: *Governing the Commons: the Evolution of Institutions for Collective Action*. Cambridge University Press (1990), 18th printing (2006)
20. Ostrom, E.: Coping with tragedies of the commons. *Annual Review of Political Science* 2, 493–535 (June 1999), http://www.cipepec.org/research/institutional_analysis/w98-24.pdf, workshop in Political Theory and Policy Analysis; Center for the Study of Institutions, Population, and Environmental Change, Indiana University, Bloomington, USA
21. Perrucci, G.P., Fitzek, F.H., Petersen, M.V.: Energy saving aspects for mobile device exploiting heterogeneous wireless networks. In: *Heterogeneous Wireless Access Networks*. Springer US (2009)
22. Perreau de Pinninck Bas, A.: *Techniques for Peer Enforcement in Multiagent Networks*. Phd thesis, Universitat Autònoma de Barcelona (2010)
23. Sabater-Mir, J.: *Trust and Reputation for agent societies*. Ph.D. thesis, Institut d’Investigació en Intel·ligència Artificial (IIIA) (2003), http://www.tesisenxarxa.net/TESIS_UAB/AVAILABLE/TDX-0123104-172828//jsm1de1.pdf
24. Simon, H.A.: Rational choice and the structure of the environment. *Psychological Review* 63(2), 129–138 (1956)
25. Wrona, K., Mähönen, P.: Analytical model of cooperation in ad hoc networks. *Telecommunication Systems* 27(2–4), 347–369 (October 2004)