Evaluating the Cost of Enforcement by Agent-based Simulation: A Wireless Mobile Grid Example

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. The subject of this paper is the cost of enforcement, to which we take a satisficing approach through the examination of marginal cost-benefit ratios. Social simulation is used to establish that less enforcement can be beneficial overall in economic terms, depending on the costs to system and/or stakeholders arising from enforcement. The results are demonstrated by means of a case study of wireless mobile grids (WMGs). In such systems the dominant strategy for economically rational users is to free-ride, i.e. to benefit from the system without contributing to it. We examine the use of enforcement agents that police the system and punish users that take but do not give. The agent-based simulation shows that a certain proportion of enforcement agents increases cooperation in WMG architectures. The novelty of the results lies in our empirical evidence for the diminishing marginal utility of enforcement agents: that is how much defection they can foreclose at what cost. We show that an increase in the number of enforcement agents does not always increase the overall benefits-cost ratio, but that with respect to satisficing, a minimum proportion of enforcement agents can be identified that yields the best results.

1 Introduction

Open systems that are sustained by the contributions, in whatever form, of the participants are fragile because of their susceptibility to free-riding. Ostrom [1] has famously documented and analysed many such scenarios, a pre-dominant feature of which is their reliance on zero or low participant turnover. Ostrom's work focuses on (semi-)closed systems in which contributions can be reinforced through mechanisms such as identity and sanctions such as ostracism. Identity is a more slippery notion in open systems, and both that and high turnover, render ineffective sanctions depending on long-term relationships and individual need for the resource created by the system. Open systems tend to rely upon other means to protect themselves and their participants, amongst which is the use of enforcement agents – a kind of police. The important distinction between such systems and Ostrom's is that behaviour is no longer self-regulating, but rather there are agents whose responsibility it is to observe behaviour, detect infractions and mete out punishment. This task then has a cost for these agents and a benefit to the system as a whole. Consequently, it is common for the system somehow to reward such agents.

As a result of the costs incurred by enforcement agents, for each such system, the questions to be answered are in particular: how many enforcement agents are needed? how much will they cost? and, how much free-riding can be tolerated? All these questions are related to the concept of "satisficing" [2] which suggests a decision-making strategy that attempts to meet criteria for adequacy, rather than optimality. We will describe the concept in more detail later in this paper.

An analytical initial estimate to answer the issues above might be possible, but begs many questions of the assumptions made about participant behaviour and effectively ignores how behaviour may change in response to system rules: will participants continue current behaviour – in the case of a rule change in an existing system – and what behaviour will participants adopt – in the case of a new system? Analytical solutions, such as game theory, also have problems in other areas such as: (1) identifying implicit but unknown effect in models, (2) enabling the testing of alternative settings and parameter sweeps, [3] which is why we opt for simulation as tool to explore the parameter space of a system that is open to free-riding and to evaluate different approaches to encouraging system sustaining behaviour.

The subject of this paper is how agent-based simulation can help answer the above kinds of questions and to illustrate the approach we take the scenario of the wireless mobile grid (WMG) [4], without prejudice as to the concept's viability. In situations where several people are interested in the same content or data (e.g. financial news, sports events,...) there are two aspects to the WMG:

1. Mobile devices download (some parts of) a resource over 3G, but opportunistically share (to acquire the whole) using peer-to-peer protocols over ad-hoc wireless networks, and
2. Infrastructure contention is lowered in densely populated situations, due to local sharing reducing the demand on 3G access.

There are individual and system benefits to each of these. The first, as well as potentially offering a speed advantage, also uses less power (the energy cost per bit for WLAN transmission is much lower than for 3G [5]) which results in longer handset standby times. The second means that network providers do not have to build infrastructure to allow for total peak demand, which may be occasional and volatile (e.g. sports events, financial districts, airports), yet can still deliver large volumes of digital content – demand for which is rising rapidly – to large numbers of users. With these benefits, WMGs contribute to a better utilization and thus sustainable handling of resources (i.e. battery capacity and networks).

The effectiveness of the WMG depends on user participation and collaboration, which raises the question of how that may be realized. The strategic user may place their own benefit above that of the collective. This matters because collaboration in a WMG comes at the cost of battery consumption. Hence, a rational user will choose to access resources without commitment from themselves. However, if a substantial proportion of users take this path, the network is unsustainable and all users are deprived of the benefits arising from cooperation [6].

However, since collaboration is essential for WMG to work and beneficial to the participants as a whole, it may be desirable to encourage it and even enforce it. Since

purely technical solutions for enforcement are frequently subverted (see [7] for example), costly to implement and difficult to replace, the WMG designers have sought to employ non-technical enforcement mechanisms in addition to potential hardware solutions. Balke [8] sets out a taxonomy of approaches for non-hardware dependent enforcement in a WMG-like environment, from which we select one particular mechanism, namely encouraging cooperation through *enforcement agents*. These enforcement agents are given the task of policing the WMG. For this purpose they participate in the system with normative power [9] and permission from the system owner to punish negative behaviour (i.e. free-riding) when detected, by imposing sanctions. Police officers are an example of enforcement agents in the real world: in contrast to "normal" citizens, they are empowered by the state to – in their function as police officers – carry out actions like arrests, etc.

This paper presents results from a social simulation developed for the purpose of analysing the effect of enforcement agents on the above outlined cooperation problem, which we situate in the context of WMGs, i.e. an example of an open-distributed system[3]. In the next section (Sec. 2), to make this paper self-contained, we summarize the key elements of WMG, how cooperation may occur, how enforcement agents help and discuss related work on mobile grids and on enforcement. Sec. 3 describes how WMGs and participating and enforcement agents are modelled in the simulation. The simulation results, their analysis and the discussion of the "satisficing" criterion are presented in Sec. 4, where the focus is strictly on the costs of enforcement related to energy consumption and cooperation benefits in the WMG.

2 Enforcement and WMGs

As noted earlier, Fitzek et al. [4] have proposed the WMG concept in order to address both energy and scalability issues arising from the high consumer up-take of edge devices utilising digital services.

The key activity in WMG is the creation of ad-hoc networks between temporarily co-located devices, in which they use short-range communication, such as wireless LAN or Bluetooth. One advantage of such communications is that, compared to 3G cellular HSDPA/LTE infrastructure communication, the energy cost per transmitted bit is much lower – approximately 1/25th or 4% of the 3G cost. For the discussion here, we assume the use of the IEEE802.11 WLAN specification, which has the highest energy saving potential and according to [5] the best potential for WMG applications.

The intrinsic weakness of the WMG proposal is that energy consumption benefits are only achievable through cooperation. In Ostrom's terms, the participants contribute resources to form a common pool [1], which all may then use to achieve a common goal, say through the downloading of part of map, video or other data fragments, which they are all interested in and which can then be combined into a whole[4]. From a utility maximisation perspective, it should be better to cooperate than act individually, but

---

[3] A preliminary and limited version of these results has been presented at a workshop. Those proceedings have not been archived.

[4] For the WMG situation situation this implies that instead of each downloading the same files, mobile phone users interested in the same data – e.g. financial news or videos from big sports

there is a risk that the cost of subscription to the WMG outweighs the gains. Thus, a (bounded) rational actor should free-ride on the WMG, except that if all actors do the same, there is no WMG.

Given that the WMG is proposed as a potential real business concept, free-riding is particularly problematic for parties trying to advance its commercial value and investing in its development. These parties are namely the telecommunications providers which, for the sake of simplicity, we assume are the same as the infrastructure providers. For these parties it is important to assess at an early stage of development of the WMG whether it is worth investing in the idea or not, by assessing:

1. Whether mechanisms can be put in place to reduce/prevent free-riding, or at least limit its impact, and
2. Whether it is possible to assess the associated costs of those mechanisms.

Enforcement agents are one such approach (others are discussed shortly), but such services are not going to be without cost themselves. Thus, it is necessary to analyse both the monetary costs (to the service providers) of employing enforcement agents, and the costs in terms of total resource consumption (in particular energy) in the system.

Having decided to examine the viability of enforcement agents in WMG, some choices are required about the form of reward and the powers they shall have within the system. It is useful to make an analogy with police officers:

1. They are paid by a state to help to ensure that the law is upheld,
2. They do not have any greater physical abilities than a "normal" citizen, but
3. They do have special rights and powers to carry out their job.

Similarly, enforcement agents participate in the WMG to observe whether free-riding is taking place, and in return receive free or reduced price phone services. We assume that the enforcement agents always report honestly, when free-riding is observed, and it is for the service provider to apply a sanction to the offending users. Sanctions could take several forms, such as limitation of the contract, contractual fines or sending battery-intensive messages to violators. This last is the option taken in this study.

2.1 Combating free-riding

There is 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 [10], who shows that in small (relatively closed) communities these "tragedies" can be overcome. However, as noted in the introduction, open systems are qualitatively different and typically require quite different solutions, for which the literature can be broken into four broad classes of solutions which can be applied separately or in combination with one another:

1. **Technological approaches:** In WMG, this means "hard-wiring and/or hard-coding restrictions" into the technological components (hardware and software) of mobile phones or into the protocols [11]. Drawbacks to technological solutions are that:

---

events – each only download part of the files and obtain the remaining parts via the more battery efficient WLAN connection, instead of 3G.

(a) adaptations at run-time can be difficult or impossible to implement,
(b) users might feel too restricted and not adopt WMG at all
(c) user modifications may render technological solutions ineffective [12, Ch.16].

2. **Formal normative concerns:** In this group works typically range from the analysis of structural relationships between system entities [13] to the discussion of operational semantics of the norms present in free-riding-like settings [14].

3. **Cognitive aspects:** Includes the adoption or internalization of norms, the role or modes of learning and adaptation that may be associated with punishment of free-riding, or the effectiveness of punishing free-riders depending on the cognitive make-up of participants [15][5].

4. **Economics-inspired approaches:** Most of the works in this category focus on punishing free-riding as a deterrent to the rational behaviour of utility-based participants (e.g. [16]). Thus, a punishment is a fine taken from the the participant's benefits. Effectiveness is usually measured against the system equilibria, that is the points where the competing (cooperation *versus* punishment) influences balance out [17]. Methodology has been either game-theoretic (see [18] for example) or experimental (including agent-based simulations) [19, 20]. Reputation mechanisms are one example here, on the grounds that a bad reputation may lead to future negative utilities (e.g. if no cooperation partners can be found due to a bad reputation) [21].

The approach adopted in this study falls into the last category above: the sanctions are fines on the free-riding agents, for which they must account in their utility considerations. In contrast to [22], we account for the costs of enforcement. These are primarily the energy costs of the enforcement agents in using their phones, but we also account for fixed overhead costs to network providers in paying each enforcement agent for its service. One paper simulating enforcement and its costs is by [23], who investigate enforcement costs in ostracism situations. For reasons already given we believe their closed world ostracism assumption is problematic for open distributed systems. Consequently, we take both the simulation approach and the proposals for which enforcement costs to consider in [23] as a starting point and adapt and extend it to account for features of the WMG domain. Agent-based simulation is the vehicle for our analysis for two reasons:

1. Simulation allows the testing of different hypotheses and parameter settings in a controlled environment, and

2. The scenario requires a heterogeneous population of phone users, whose proportions and individual behaviours can be varied.

The users are mobile, have limited information about the system and are assumed to make bounded rational decisions on the basis of their goals and the environment they operate in. This highly complex setting is inappropriate for analytical modelling, so an empirical approach remains the best way to test our hypotheses about the cost and effectiveness of enforcement.

---

[5] As such, reputation mechanisms have a place here, if the participants are cognitive agents that are affected by poor reputation. We discuss the economic perspective of reputation under the next heading.

3 The WMG simulation

To describe the experimental set-up of our agent-based simulation, we explain the three component models of which it is comprised:

1. The technical model of the WMG,
2. The model of the actors in the system represented as agents,
3. The cost model for the introduction of enforcement agents

Subsequently, we present the simulation environment, its controlling parameters and the hypotheses to test.

3.1 The Technical Model

The design of the technical model is based on a few assumptions commonly known together as the "Flat Earth" model [24]. This is widely used and accepted as 'best practice' in communication simulations, hence we adopt it here. The relevant assumptions are symmetry (if node A can hear node B, B can hear A), perfect signal strength, no external interference, no obstacles and limited communication range. Furthermore, we assume that all participants in the system have a mobile phone with the same technical specifications, that is having equal battery consumption rates for equal actions performed. This assumption is supported by [5] which states that for representative phones with features for WMG communication, the differences between different mobile phones are only marginal. We use his measurements of mobile phone battery consumption for sending and receiving on 3G and WLAN respectively.

3.2 The Agent and Interaction Model

We distinguish two types of system participants: "normal" users and enforcement agents. Both types move randomly and at any given point of time can interact with the agents that are in their communication radius, as determined by the operational range of the WLAN connection.

For the normal agents in the system, we assume bounded rationality and (personal) utility maximization. The latter means that when deciding what to do, the agents consider the utility of the available actions and choose the one with the highest utility. Utilities may vary from agent to agent. An agent could for example assign cooperation a high utility, whereas another might assign a higher utility to the reduction of their battery consumption. Bounded rationality specifies that the agents do not necessarily know the whole system and thus cannot base their decision on complete (system) knowledge. That is why they can only optimize their local (current location bound) utility, which may be different from the global one. The procedure for the decision-making process for normal agents and their utility considerations is described in Figure 1, which we will now explain in more detail: The normal agent decision-making procedure is:

1. At simulation start, normal agents are randomly assigned the task of downloading a specific file. On receiving this task, the agents can decide on how to obtain this file, that is either by downloading it themselves or by trying to find a collaboration

Fig. 1. The Agent Decision Making Process. The numbers in the decision nodes correspond with the item numbers in the decision-making procedure described in section 3.2.

partner with which to obtain the file jointly. The former option is of particular interest if there are few neighbouring agents (i.e. potential collaboration partners) in the vicinity or if the file is small, as both situations could result in low cooperation gains as well as relatively high costs for finding cooperation partners.

2. If the number of neighbours and the file size are sufficiently high[6], the agent broadcasts cooperation request that specifies the file the agent wants to download and its request for cooperation. Requests for collaboration do not need to be sent if the agent receives a matching cooperation request from another agent, to which it can respond.

3. An agents seeking cooperation partners waits a fixed amount of time for responses. Having received enough positive responses, the agent checks who has responded and decides whether it wants to collaborate with the other agents. Reasons for not wanting to cooperate could be bad past behaviour for example.

4. Having agreed to join a cooperation group, the agent has two tasks. First of all it has to decide whether to download its promised share from the base-station and more importantly, having downloaded its share, to decide whether to cooperate further, that is to send the share to its cooperation partners, or to defect.

5. The cooperation decision is based on the utilities the agent assigns to cooperation and defection including is assumptions about the behaviour of other agents as well as the potential chances of detection and of being fined.

6. Having made its decision, the last step is to wait and see whether the cooperation partners send the promised shares. If shares are missing, the agent must repeat the decision process and decide whether to download the missing shares itself or find new cooperation partners.

In contrast to the normal agents, the enforcement agents are not given a download task, but only become active when they receive a collaboration proposal from another agent. If an enforcement agent is not already engaged in an interaction, it accepts the proposal and always sends its share. The energy consumption costs that the enforcement agents incur for sending their share, are added to the total energy consumption in the WMG in order to account for the additional energy that the enforcement mechanism consumes. After the deadline for cooperation has passed, the enforcement agent identifies the cooperation group members that did not cooperate. These agents are reported and a sanction applied to them. It is important to note with respect to defection from the agreed cooperation, that according to the current rules, an agent may not be punished more than once for a defection. Hence, if several enforcement agents detect the same violation, the violator is only punished once.

3.3 The Cost Model

Having described the WMG model and the agent (decision) model, the last aspect we consider is the analysis of the implications of the employment of enforcement agents.

---

[6] This value is agent-specific. It is based on an agent's attitude to risk and preference for cooperation.

In view of the potential commercial interest in WMGs, from the large number of possible factors to evaluate with the simulation, we concentrate on two linked to business: financial implications – what are the financial overheads associated with the mechanism, as well as the potential benefits – and overall battery consumption – potential reduced consumption and extended handset cycles make this attractive for user. To do so, the following two assumptions are made:

1. The one universal mobile phone provider offers its customers the option to join the WMG or not. If they join and the cooperation is successful, they can reduce their battery consumption according to the specifications given earlier. Customers can opt not to participate – that is never accept a collaboration request – without penalty. However if a user does accept a collaboration request, but fails to contribute, he/she runs the risk of being penalised.

2. The costs of the enforcement agents are two-fold: firstly, as mentioned earlier, enforcement produces additional battery consumption that needs to be accounted for. Secondly, in this paper, we assume that in return for their support in controlling the wireless mobile grid and as incentive to report truthfully, they can use their mobile phones free of charge and thus can save an amount c_{eAg} which is the average level of a mobile phone bill. In consequence, for each enforcement agent employed, the infrastructure provider loses c_{eAg} of potential revenue. Here, we account for both types of costs, starting with the energy consumption and afterwards extend it to the reward costs for enforcement agents. This includes a discussion of the significance of the level of c_{eAg} in Section 4, where we analyse the results of the simulation experiment. As also mentioned earlier, for the sake of simplicity and because they benefit from honesty by saving c_{eAg} we assume that the enforcement agents always try to fulfil their duty. It can however happen that because agents move out of range of one other, they sanction agents that act honestly, that is produce statistical false-positives. We are able to check for these and consider them in the analysis of the success of the enforcement mechanism.

3.4 Simulation Parameters

The implementation of the above components, namely the agents and the technical and the financial model, is realized using the Jason simulation platform [25]. The agent reasoning is encoded in AgentSpeak, and all remaining parts of the simulation are programmed in Java. Table 1 gives an overview of the simulation configuration parameters used in the experiments.

The simulation parameters specify the number of the normal agents, the duration of the simulation (i.e. the number of tasks – the number of files to obtain – per agent in the simulation). We use a random walk to determine where an agent is located at each cycle and thus which neighbours and potential cooperation partners it has. The communication radius is based on WLAN specifications and is the same for all agents (including the enforcement ones).

The purpose of the simulation is to test not only the impact of enforcement agents, but also the sensitivity of the cooperation problem to the number of such agents in proportion to the normal population. Furthermore we investigate the effect of the level of

Table 1. Simulation Variables

Name	Simulation Parameter
# Normal Agents	400
Enforcement Agents as % of Normal Agents	0%, 1%, 2%, 3%, 4%, 5%
Fine percentage	100%, 200%, 300%, 400%, 500%, 600%, 700%, 800%, 900%, 1000%
# Consecutive Tasks	50
Movement Pattern	Random Walk

the fine issued by the enforcement agent in case of a violation. The primary performance indicator is taken to be the total battery consumption across the entire population, from which we exclude the effect of fines following a violation, to be able concentrate on the energy required for downloading and sharing files (as well as policing the system) in our analysis. We identify a measure we call the *average energy consumption (AEC) ratio*, which is the ratio of the actual battery costs of the interactions that take place, over the theoretical costs of downloading if each participant downloads everything themselves. As more sharing occurs, this ratio tends to 0, conversely, less sharing and the ratio tends to 1 (or theoretically even above 1, if the defection rate is high). Consequently, the simulation experiments test the following hypotheses:

Hypothesis 1: The employment of enforcement agents increases cooperation and decreases the AEC ratio.

Hypothesis 2: An increase in the percentage of enforcement agents results in a reduction of the AEC ratio.

Hypothesis 3: An increase in the fine value results in an increase in cooperation and a reduction of the AEC ratio.

In order to test these hypotheses we run factorial experiments, in which each parameter ranges over a discrete set of possible values and the set of experiments cover all the combinations of values across the factors. Each experiment consists of 50 simulation runs for each of the parameter combinations in Table 1. For each run, each agent has the task of downloading a file – and hence making cooperation decisions – 50 times. Each time, the file assigned to the agent is randomly chosen from a pool of files. The results of these runs are then analysed using Matlab. We use ANalysis Of VAriance (ANOVA) to test the significance relationship between the parameters in the simulation (independent variables) and the number and ratio of defections and the AEC ratio (dependent variables). This informs us whether specific input parameters (or parameter combinations) influence the outputs[7]. As a post-hoc test to ANOVA, we also use Tukey's test, which identifies the impact of specific variables on the overall result.

---

[7] We performed the Shapiro-Wilk test as well as Levene's test to ensure the applicability of the ANOVA.
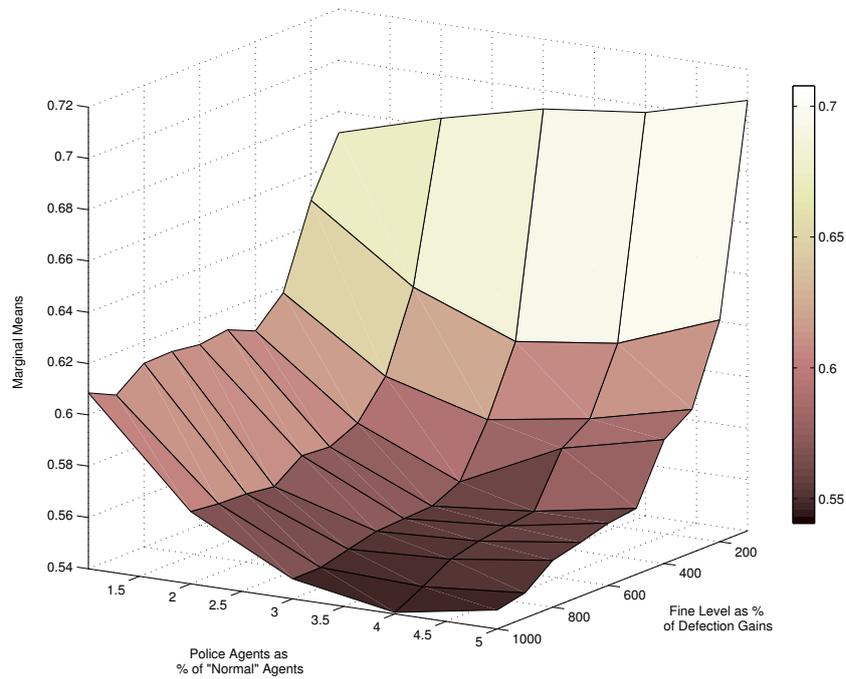
Fig. 2. Marginal Means for Comparing Simulation Experiments with different Percentages of Enforcement Agents for different Fine Levels. In a design with several factors (e.g. battery capacity, police agent percentage and fine level), the marginal means for one factor are the means for that factor averaged across all levels of the other factors.

4 Results & Evaluation

Following the presentation of our hypotheses and the experimental set-up, we now turn to the analysis of the simulation results in order to confirm or refute the proposed hypotheses.

We start with hypothesis one which states that the introduction of enforcement agents increases cooperation and reduces the overall AEC in the system. To test this hypothesis we look at differences in the AEC rates in the simulation experiments with and without enforcement agents and compare these for the different fine values. With the help of ANOVA, we test whether the null hypothesis that the enforcement agents account for no difference can be rejected or not. We performed this test for all fine levels and get significant p-values (< 0.0001) for all of them, that is, we can reject the null hypothesis and conclude that the utilisation of enforcement agents results in a difference in the AEC rates. This is true even for low fine levels and indicates the deterrent effect of fining. We use one-tailed t-tests to establish the nature of the difference between having

enforcement agents and having none and confirm that lower AEC ratios are dominant in the simulation experiments with enforcement agents. This confirms Hypothesis 1. However, when performing the ANOVA, high error rates for all fine levels are detected. This typically indicates that a difference exists between the rates of enforcement agents and fine levels 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 agent percentage and fine level contributes to this difference, we perform Tukey's test as a post-hoc analysis. Figure 2 shows the results of a Tukey's test analysing this effect. It displays AEC ratio marginal means for the different simulation parameter setting combinations. The lower these marginal means (i.e. the darker the display colour), the less energy overall is consumed in the system. The figure does not include the results of the baseline case simulation experiments without any enforcement agents as these distort the graphical representation of the remaining data-points. For reference, the marginal means for settings without enforcement agents have values of approx. 0.7843.

Looking at Figure 2 in more detail: two effects can be observed which are important for our analysis:

1. An increase in the fine leads, in most cases, to lower AEC ratios
2. An increase in the percentage of enforcement agents does not always result in a reduction of the AEC ratio, implying Hypothesis 2 is incorrect.

Looking at the latter effect in more detail first, Figure 2 indicates that for enforcement agent levels of 0% to 3%, Hypothesis 1 seems true: a steady (although not linear) decrease in the energy consumption ratios can be seen for all fine levels (downward direction of the surface with an increase of enforcement agents from 1% to 3%). However this picture changes when higher number of enforcement agents are employed, where the AEC ratio increases instead of decreases and the surface in the plot goes upwards. This initially surprising effect, can be explained by economic principles underlying the simulation experiments, as enforcement agents have associated costs. Despite the reduction of battery usage when more enforcement agents are employed, this additional reduction – and the detection of violations – does come at a price, which must be summed across all the enforcement agents. In contrast, the detection of violations follows the economic principle of the law of diminishing marginal utility (also known as "Gossen's First Law"). This law states that for any good or service, its marginal utility decreases as the quantity of the good or service increases, *ceteris paribus*. In terms of our simulation this implies that: with increasing numbers of enforcement agents, the number of additional detections ($\Delta(D)$) decreases with each additional enforcement agent ($\Delta(x)$). This can be explained by:

1. The larger the number of enforcement agents, the higher the probability that several enforcement agents are in the same location and observe the same violation. We pointed out earlier that – in accordance with general legal principles – an agent can only be held accountable once. Thus, an agent can only be punished once regardless of how many enforcement agents observe the violation. As a result, additional observations of the same event do not produce any additional benefit. Nevertheless, the costs for the enforcement agents need to be paid: that is, their battery consumption still counts towards the total battery consumption across the system.

2. In the case of successful enforcement, that is a high number of detections, rational agents will attempt to cooperate in order to avoid a possible detection of violations and consequent fine. So, the total number of violations decreases, automatically reducing the number of potential detections, resulting in the same enforcement agent cost problem.

As a consequence, instead of aiming for absolute enforcement, economic theory suggests using "satisficing" [2], that is *minimizing $\Delta(D)/\Delta(x)$*, which can be determined by means of the factorial experiments, and the best result approximated. Additionally, a simulation can also be used for testing thresholds, for example, how many enforcement agents are needed on average to detect 90% of all violations. Thus, quality of service requirements can be stated and the associated costs determined *a priori*.

In our experiments – comparing 3%, 4% and 5% enforcement agents – the lower percentage of enforcement agents performs better for satisficing cooperation with respect to the AEC ratio. Conversely, for between 0% and 3% enforcement agents, the gains made by additional units of enforcement agents *are* beneficial, suggesting there is a minimum in this region. Thus, although not optimal with regard to the detection of violations (3% enforcement agents will detect less than 5%) the costs associated with them, that is 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. However, as pointed out earlier, the WMG is a real world business case study and additional costs/losses might be incurred. When first describing the idea of enforcement agents, we explained that one incentive for them to help control the system could be financial benefits, such as free mobile phone contracts. This incentive can be quantified. Currently, mobile telecommunication providers have an average monthly revenue per user of approximately $c_{eAg} = €\ 15$ [26], which is mainly generated through contracts. For 3% enforcement agents, for a provider like T-Mobile with 35,403,000 customers in Europe in 2011 [27, p.90], this implies a monthly loss of revenue of approximately €\ 15.9 million. This cost needs to be weighed against the benefits of the WMG, such as energy savings or possibly higher subscriptions numbers if users embrace the WMG idea and potential capital savings on network infrastructure. Based on these figures, 3% enforcement agents seems a rather high price, and a lower number of enforcement agents or a less generous incentive for their service (e.g. only covering the marginal costs of their contracts) might seem more affordable.

One other means to decrease the AEC ratio is to increase the fine level and thereby indirectly decrease the number of defections. In Figure 2, we can observe that the increase in the fine level indeed decreases the mean values (confirming Hypothesis 3), however the decrease has a diminishing nature. This suggests a saturation effect with respect to the fine level, that is the higher the fine gets, the lower the additional gains from it. Consequently, we observe that high fine levels do not yield high benefits and can even increase battery consumption, as energy used by the enforcement agents does not result in the detection of violations. On a social level, they could also result in acceptance problems by the wireless grid users, which in turn is counter-productive if infrastructure providers want to promote WMGs as part of their business concept.

The feared problem of false-positives in our simulation (e.g. false accusations due to agents moving out of each other reception range) did not have a major impact in the

simulation experiments. Over all experiments we had 0.007% false positive rate, i.e. 7 in 100,000 users were wrongly accused.

5 Conclusion & Future Work

We have presented an analysis of the costs of enforcement, depending on two factors: the amount of enforcement and the fine level imposed in case of sanction.

We use our analysis to examine the facilitation of cooperation in WMGs – a kind of large-scale open distributed system. For the WMG case study we are able to confirm two out of our three hypotheses, namely that enforcement agents help to reduce the AEC ratio and that an increase in the fine level is aligned with this effect. We furthermore establish, that the increase in both the number of enforcement agents and the fine level have diminishing returns, which in the case of the number of enforcement agents can even lead to negative battery consumption results. Hence, in combining the two concepts it is important to get the right balance between the deterrent effect of the enforcement mechanisms and the costs associated with it. Although these results are based on a case-study specific agent-based simulation, our experiments clearly demonstrate that the costs associated with enforcement are important to consider in any enforcement mechanism, as they can significantly alter its benefit. Clearly each scenario has its own key parameters, so while the results of this study cannot be applicable in another domain, we believe the approach offers useful methodological lessons for the modelling and evaluation of mechanisms for the maintenance of open system common pool problems in general.

As part of future work, we intend to investigate further the effect of the movement of agents, by implementing more realistic movement patterns than random walk. We also plan to explore a wider range of evaluation parameters. On a different level, the introduction of coordination between the enforcement agents also seems sensible. Currently, the individual enforcement agents act independently. Despite this lack of coordination, we only recorded a very few occasions in which two enforcement agents observed the same interaction, which is not surprising given the relatively small percentages of these agents in the population. Nonetheless, a natural human assumption, which may be overturned by the results, would be that for more effective enforcement, a limited level of coordination between these agents could be useful.

Finally, our case-study in based on direct reciprocation, in which collaboration always focuses on an immediate quid pro quo. This neglects the idea that the contribution of an agent to the system and its benefit from the system might have a time discrepancy, such as in routing scenarios. Accounting for a lack of direct interaction is therefore also a part of future work.

References

1. Ostrom, E.: Understanding Institutional Diversity. Princeton University Press (2005)
2. Simon, H.A.: Rational choice and the structure of the environment. Psychological Review **63**(2) (1956) 129–138

3. Conte, R., Gilbert, N.: Computer simualtion for social theory. In: Artificial Societies: The Computer Simulation of Social Life. Routledge (1995)

4. 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. Springer (2007) 31–59

5. 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)

6. Wrona, K., Mähönen, P.: Analytical model of cooperation in ad hoc networks. Telecommunication Systems 27(2–4) (October 2004) 347–369

7. Leibowitz, N., Ripeanu, M., Wierzbicki, A.: Deconstructing the kazaa network. In: Proceedings of the Third IEEE Workshop on Internet Applications, IEEE Computer Society (2003)

8. Balke, T.: A taxonomy for ensuring institutional compliance in utility computing. In Boella, G., Noriega, P., Pigozzi, G., Verhagen, H., eds.: Normative Multi-Agent Systems. Number 09121 in Dagstuhl Seminar Proceedings, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Germany (2009)

9. Jones, A.J.I., Sergot, M.J.: A formal characterisation of institutionalised power. Logic Journal of the IGPL 4(3) (1996) 427–443

10. Ostrom, E.: Coping with tragedies of the commons. Annual Review of Political Science 2 (June 1999) 493–535 Workshop in Political Theory and Policy Analysis; Center for the Study of Institutions, Population, and Environmental Change, Indiana University, Bloomington, USA.

11. Chen, W., Guha, R.K., Kwon, T.J., Lee, J., Hsu, Y.Y.: A survey and challenges in routing and data dissemination in vehicular ad hoc networks. Wireless Communications & Mobile Computing 11(7) (2011) 787–795

12. Schneier, B.: Liars and Outliers. Wiley (2012) ISBN13:978-1-118-14330-8.

13. Boella, G., Pigozzi, G., van der Torre, L.: Normative framework for normative system change. In: AAMAS (1). (2009) 169–176

14. Grossi, D., Aldewereld, H.M., Dignum, F.: Ubi lex, ibi poena: Designing norm enforcement in e-institutions. In Noriega, P., Vázquez-Salceda, J., Boella, G., Boissier, O., Dignum, M., Fornara, N., Matson, E., eds.: COIN II, Springer (2007) 101–114

15. Conte, R., Andrighetto, G., Campenni, M.: Internalizing norms. a cognitive model of (social) norms' internalization. The International Journal of Agent Technologies and Systems (IJATS) 2(1) (2010) 63–73

16. 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)

17. Becker, G.S.: Crime and punishment: An economic approach. The Journal of Political Economy 76(2) (1968) 169–217

18. Coleman, J.S.: Foundations of Social Theory. Belknap Press (August 1998)

19. Gurerk, O., Irlenbusch, B., Rockenbach, B.: The competitive advantage of sanctioning institutions. Science 312(5770) (April 2006) 108–111

20. Perreau de Pinninck Bas, A.: Techniques for Peer Enforcement in Multiagent Networks. PhD thesis, Universitat Autónoma de Barcelona, Spain (2010)

21. König, S., Balke, T., Quattrociocchi, W., Paolucci, M., Eymann, T.: On the effects of reputation in the internet of services. In: Proceedings of the 1st Int. Conference on Reputation (ICORE 2009). (March 2009) 200–214

22. Andrighetto, G., Villatoro, D.: Beyond the carrot and stick approach to enforcement: An agent-based model. In: European Conference on Cognitive Science. (2011)

23. El Mouden, C., West, S.A., Gardner, A.: The enforcement of cooperation by policing. Evolution **64**(7) (2010) 2139–2152
24. 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, ACM (2004) 78–82
25. Bordini, R.H., Hübner, J.F., Wooldridge, M.: Programming Multi-Agent Systems in AgentSpeak using Jason. Wiley Series in Agent Technology. John Wiley & Sons (2007)
26. ie Market Research: 3q10 germany mobile operator forecast, 2010 – 2014. `https://www.iemarketresearch.com/Members/Reports/3Q10-Germany-Mobile-Operator-Forecast-2010--2014-Germany-to-have-118-million-mobile-subscribers-in-2014-with-market-share-of-T-Mobile-declining-to-34--RID1590-1.aspx` (2010)
27. T-Mobile: Geschäftsbericht (annual report). `http://www.telekom.com/static/-/103442/10/120223-ar2011-pdf-si` (2011)