



Citation for published version:

Balke, T, De Vos, M & Padget, J 2011, 'Analysing energy-incentivized cooperation in next generation mobile networks using normative frameworks and an agent-based simulation', *Future Generation Computer Systems*, vol. 27, no. 8, pp. 1092-1102. <https://doi.org/10.1016/j.future.2011.04.006>

DOI:

[10.1016/j.future.2011.04.006](https://doi.org/10.1016/j.future.2011.04.006)

Publication date:

2011

Document Version

Peer reviewed version

[Link to publication](#)

University of Bath

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.

Analysing Energy-Incentivized Cooperation in Next Generation Mobile Networks using Normative Frameworks and an Agent-Based Simulation

Tina Balke^a, Marina De Vos^b, Julian Padget^b

^a*University of Bayreuth
Chair for Information Systems Management
Bayreuth, Germany*
^b*University of Bath
Department of Computer Science
Bath, UK*

Abstract

Wireless mobile grids (WMG) have been proposed as the next generation mobile networks in order to address the energy issues arising for the next generation of mobile phones. WMG are based on the notion that local communication is more energy-efficient than the standard 3G communication. Despite their energy advantages, they create a social dilemma as it is advantageous for rational users to benefit from the energy savings without any contributing to the cooperation, as every commitment has its price. This paper proposes the use of a normative framework modelling technique and multi-agent simulations to support the early and rapid development of prototype systems to analyse solutions for solving the collaboration dilemma in WMG. Both tools allow for the capture of essential features of interactions between components in open architectures, therefore assisting in concept refinement, identification of actors, policy exploration and the feasibility assessment of new systems. With the help of these models we show how it is possible to quantify energy consumption, explore management policies and evaluate individual utility functions, all of which act as drivers in helping realise the WMG concept.

Keywords: Wireless Mobile Grids, Energy reduction, Simulation, Normative reasoning, Reciprocity, Multi-agent systems

Email addresses: tina.balke@uni-bayreuth.de (Tina Balke),
mdv@cs.bath.ac.uk (Marina De Vos), jap@cs.bath.ac.uk (Julian Padget)

1. Introduction

The development and deployment costs for new technology and infrastructure in mobile communication are extremely high both from a time and financial perspective. Sufficiently rich simulation environments can significantly reduce the costs. Thanks to their flexibility and rapid prototyping approach, models and strategies can easily be tested and evaluated. This paper focusses on cost-efficient techniques for modelling and the evaluation of strategies for reducing battery consumption in next-generation mobile networks (NGMNs)¹. The novelty of our approach lies in the combination of a formal model—called a normative framework—that permits design-time verification of model properties, with a bottom-up agent-based simulation, that offers empirical runtime validation of system behaviour. In consequence, we be confident in the accuracy of the formal model, while at the same time being able to look for both gross system properties and emergent properties that would not otherwise be apparent in an design-time model.

The particular NGMN scenario we are interested in are wireless mobile grids (WMG) that have been proposed to address the energy issues in NGMN. At present, the deployment of third generation (3G) of mobile network systems is in progress, but a quite different next generation network (called Fourth Generation or 4G) is under development that is intended to cause a paradigm shift in the cooperation architecture of wireless communication Katz and Fitzek (2006), with the main to drastically reduce the battery consumption.

Batteries have fixed capacity that puts limits on the operational time for a device in one charge cycle. The increasing sophistication of mobile phones and their evolution into smart phones offering Internet access, imaging (still and video), audio and access to new services, has had a significant impact on power consump-

¹We understand the term NGMN as defined by the NGMN Alliance as distinct notation for the next generation of mobile wireless networks. The term “NGMN” is technology-agnostic and does not refer to one specific technology or technical standard, but to a range of technologies and standards with certain minimum characteristics, of which the most prominent ones are:

- The support of data rates up to 100 Mbit/s in the downlink (i.e. from the network to the end-user terminal) and up to 50 Mbit/s in the uplink (i.e. from the end-user terminal to the network) within a 20 MHz channel bandwidth; and
- a low end-to-end latency (round trip time for data packets) of less than 30 ms (Alliance, 2006).

tion, leading to shorter stand-by times, as well as the problem of rising battery temperature unless there is active cooling (Perrucci et al., 2009).

Fitzek and Katz (2007) have proposed a way around some of these issues with the concept of a ‘WMG’, in which users share resources in a peer-to-peer fashion that uses less power but this requires a difficult to obtain collaboration between the users. The social dilemma thereby is that the network users can show strategic behaviour. Thus the main problem in WMG is that commitment in these networks comes at a cost in the form of battery consumption, etc. As a consequence, rational users would prefer to access the resources without any own commitment. However, if a substantial number of users would follow this selfish strategy, the network itself would be at stake, depriving all users from the benefits (Wrona and Mähönen, 2006).

In Balke et al. (2011) we reported on a feasibility study into how and whether normative frameworks can assist in evaluating the functionality of NGMN and presented a design-time (also called off-line) model focussing on enforcing cooperation. Starting from this idea, the contribution of this paper is two-fold: (i) We build a normative design-time model of the wireless grid cooperation problem in order to provide a foundational model to quantify the potential WMG benefits if enforcement can take place. (ii) The second contribution is the construction of a complementary empirical agent-based simulation, where in contrast to the design-time model, we model the (normative) behaviour of individual actors—in this case, the handsets (or their users: the distinction is unimportant)—participating in the WMG. Through such a simulation, we can collect empirical evidence that is the accumulation of the local decision processes of significant populations of handsets, to see how individual utility functions contribute to or hinder the realization of the WMG vision. After the construction of the two models (design and runtime), in a second step we show how the design-time and the runtime model (within the simulation) work together and how their combination helps to analyse the cooperation problem.

The remainder of the paper is structured as follows: in the next section (Section 2) we present a detailed description of the NGMN scenario that we use as a case-study in this paper (i.e. the idea of WMG), an analysis of the energy costs of the different networks that have motivated the idea and examine the reciprocity problem in the framework of the classical game theory Prisoner’s Dilemma model, concluding that (without enforcement mechanism) non-collaboration is the preferred course of action. However, that is precisely *not* the outcome we seek. For the collaboration to succeed, individuals must have suitable incentives not to defect, and it is the purpose of the normative framework to provide the context in



Figure 1: WMG Communication Architecture

which those incentives are attractive. Thus, in Section 3, we briefly outline the formal modelling framework we have developed and its computational realization in Answer Set Programming, before moving on (Section 4) to describe a *design-time* WMG model for an idealized scenario in which some actors download some parts of some digital content individually, but also share the downloaded parts with each other. This enables us to quantify, in terms of battery consumption the benefits of using a WMG. In Section 5, we describe the *runtime* agent-based simulation, where we explore the same scenario from the perspective of individual decision-makers and can collect data about overall system behaviour, in order to contrast it with that of the normative model. In conclusion, in Section 6, we discuss how these two views on modelling potential future technological developments complement one another and consider how they may be connected in future so that high-level system objectives may directly influence low-level decision making in distributed, adaptive systems. Please note that, at variance with standard practice, there is no single related work section. Instead, because of the way in which our contributions combines research from several areas, we discuss and cite key papers in the relevant sections.

2. The Reciprocity Problem in Wireless Mobile Grids

2.1. The Wireless Mobile Grid Architecture

As described in section 1, to overcome the energy problems of NGMN, Fitzek and Katz (2007) proposed the establishment of so-called WMG as shown in Figure 1. In WMG, ubiquitous mobile devices with potentially different capabilities are expected to build ad-hoc connections, to cooperate and to share their limited resources for the benefit of the community. The cooperation between devices is realised through the use of a short range communication link, such as WLAN or

Bluetooth. Compared to the standard cellular HSDPA communication via a base-station, the advantage of short-range communication is much higher bandwidth at the cost of much less power, as shown Section 2.2 (Perrucci et al., 2009). Here we focus on the IEEE802.11 WLAN specification, that permits mobile devices to communicate directly with each other and, according to Perrucci et al. (2009), has the highest energy saving potential².

For a better understanding of the WMG proposal, we briefly present a scenario for use in the rest of the paper. The scenario is London’s financial district, which is highly interesting from an infrastructure provider’s point of view, because of the high number of potential customers as well as the problems arising from the high density of mobile phone users. In consequence, the network may easily become overloaded and quality of service deteriorate. The reason for this is straightforward: we may assume some network users want to download video-streams such as financial news from a single base station, that uses the conventional multicast technique, whereby the bandwidth of the base station is divided into several sub-slots (‘channels’), that are sent out sequentially within one time frame. Thus – up to a technology-defined maximum – each mobile phone is assigned one slot. As the total bandwidth of a base station is fixed, the more mobile phone users are assigned a slot, the smaller the bandwidth gets that can be allotted to a single channel. As a result download times increase, leading both to higher battery consumption as well as lower quality of the streaming service.

In contrast to the normal ‘non-cooperative’ scenario, where a single mobile phone user would need to receive all sub-streams over the cellular link, resulting in the problems identified above, cooperation in the form of a WMG, enables users to share the task by receiving a subset of the multicast channels over the cellular link from and acquiring the remaining parts over the short range link.

2.2. *The Energy Advantage in IEEE802.11*

To understand the IEEE802.11 WLAN WMG scenario and its energy implications better, this section details various technical aspects of the communication costs. We can express energy consumption E in terms of two factors: the power P consumed per connection type and the time t taken for transmission:

$$Energy = Power * Time [Joules] \tag{1}$$

²In contrast to 802.11 WLAN, Bluetooth does not permit direct communication between individual mobile devices, but rather the concept of a master device that is connected to a maximum number of seven slave devices. The slave devices can only communicate via their master, hence more data exchange is needed.

Table 1: Power Level and Data Rate for Cellular - 100 byte

state	power value [W]	data rate [Mbps]
receiving	1,314	0,193
idle	0,661	-

Table 2: Power Level and Data Rate for WLAN Broadcast - 1000 byte

state	power value [W]	data rate [Mbps]
sending	1,629	5,623
receiving @ 3m	1,375	5,379
receiving @ 30 m	1,213	5,115
idle @ 3m	0,979	-
idle @ 30m	0,952	-

The total energy consumption is the energy consumed over the cellular 3G connection (E_{3G}) plus that over the short link (i.e. WLAN) connection (E_{WLAN}). In case of no cooperation the latter costs 0, i.e. it is assumed that the WLAN connection is turned off and the user has to stream the complete video using the 3G connection. In case of WMG, it is assumed that both connections (WLAN and 3G) are turned on and the devices help one another in a peer-to-peer-like fashion. We use A to denote the set of users and A_{Coop} to denote those that are cooperating. Thus given $|A_{Coop}|$ cooperating users, each only needs to stream a part of the total video via the base station (i.e. $\frac{1}{|A_{Coop}|}$ in an ideal scenario) and obtain the remainder from the cooperation partners via the short link. Thus, total energy consumption in the cooperative case (E_{Coop}) comprises³ that for:

- streaming part of the video from the base station using the 3G link: $E_{3G,rx}$,
- receiving the remaining video chunks on the WLAN connection: $E_{WLAN,rx}$,
- transmitting the user's own chunks to the other participants via the WLAN connection: $E_{WLAN,tx}$, and
- idling, that is when not transmitting or receiving anything but waiting for the next interaction on either connections: $E_{WLAN,i} + E_{3G,i}$

³Of course in order to receive at optimal results, the users need to negotiate which user is downloading which chunk and each user only asks for that specific chunk from the base station. For a simpler presentation and analysis this issue is neglected here, but we account for it in the latter part of the paper.

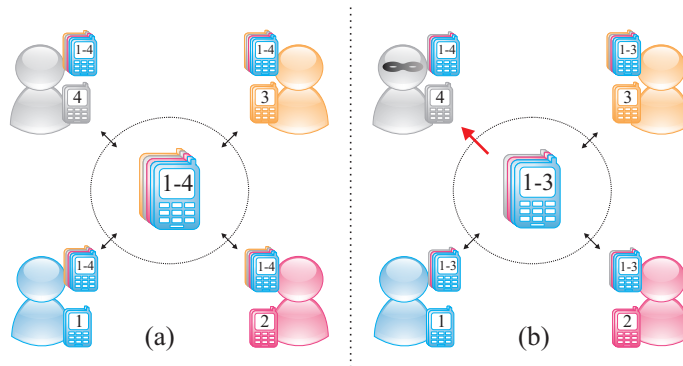


Figure 2: The Reciprocity Problem in WMG

By substituting the respective $P * t$ -values into Equation 1, one can analyse the power consumption as well as the transmission times of the scenario in the cooperative and non-cooperative case in detail. Representative power and time values for the transmission in different states using 3G and WLAN connection are reported in Perrucci et al. (2009, p. D10)⁴. We reproduce their results in Tables 1 and 2. These figures are for a Nokia N95 running Symbian OS v9.2. They show that although the power needed for the WLAN and the 3G state are about the same, the data rate for the 3G link (0.193 Mbit/s for the receiving state) is significantly lower than that of WLAN (5.115 Mbit/s, receiving state, 30m distance) leading to a significantly worse energy per bit ratio for the 3G link⁵.

This suggests that the cooperation scenario has a significant *potential* advantage in energy consumption, compared to the conventional cellular communication architecture, especially if the number of cooperating mobile phones is high and a large proportion of the data transmission can be done via the short-link connection.

2.3. The Reciprocity Problem in Wireless Mobile Grids

Although the WMG scenario suggests potential advantages with regard to battery consumption, it also has the intrinsic weakness of any distributed cooperative architecture: it relies on cooperation to succeed. We depict the idea of cooperation in WMG in Figure 2(a), and describe it as follows:

⁴According to the authors of the paper, these data are representative for large amounts of data transmitted.

⁵The energy consumed in the idle states is not measured in (Perrucci et al., 2009).

- The participants volunteer their resources, forming a common pool which can be used by all, in order to achieve a common goal, such as file streaming. The utility that users can obtain from the pooled resources is higher than they achievable individually. However, the problem is that commitment comes at a cost, in the form of battery consumption for sending file chunks, i.e. $E_{WLAN,tx}$.
- In consequence, rational users would prefer to access the resources without any commitment of their own. Thus, as shown in Figure 2(b), the grey user in the top left corner (with mask) can enjoy the advantages of the common pool without committing anything itself, hence cheating on the three other users.

However, if a substantial proportion of users follow this selfish strategy, the network itself is at stake, depriving all users of the benefits (Wrona and Mähönen, 2006). The reason for this is straightforward: network users can have strategic behaviour and do not necessarily obediently cooperate in making their resources available, without the prospect of reward for good behaviour. Unreciprocated, there is no inherent value to cooperation for a user. A lone cooperating user draws no benefit, even if the rest of the network does. Guaranteed cost paired with uncertainty or even lack of any resulting benefit does not induce cooperation in a rational, utility-maximising user. Without any further incentives, rational users therefore would not cooperate in such an environment and all be worse off than if they cooperated (Axelrod, 1981; Ostrom, 1999).

Let us consider the problem from a game-theoretic point of view: the individuals are caught in a prisoners dilemma (PD) as shown in Figure 3.

		Cooperation partner n	
		cooperation	no cooperation
User	cooperation	$E_{no\ coop} - E_{coop}$	$E_{no\ coop} - (E_{coop} - E_{WLAN,tx})$
	no cooperation	$E_{no\ coop} - E_{coop}$	$-E_{WLAN,tx}$

Figure 3: The Free Riding Problem as cardinal PD

		Cooperation partner n	
		cooperation	no cooperation
User	cooperation	2	1
	no cooperation	4	3

Figure 4: The Free Riding Problem as ordinal PD

Figure 3 shows the relative payoffs for a WMG user and a potential cooperation partner. The payoffs are shown in terms of energy savings for each transaction

partner in comparison to the option of a participant downloading everything itself. We distinguish three cases:

- In case of full cooperation the energy saving potential is equivalent to $E_{no\ coop} - E_{coop}$ (i.e. simply the energy gained by cooperating in contrast to doing everything oneself) for each partner⁶.
- In case where one partner defects, it gains an additional payoff $E_{WLAN,tx}$, because it saves energy by not sending its file chunks to its partner. The other partner does not receive the file chunks and has to download the remaining ones itself. Consequently, it downloads everything itself (i.e. has no energy saving with regard to the download) and in addition has additional energy costs for sending its file chunks.
- The last case shows the payoffs for both partners defecting. In this case they both have to download everything themselves leading to a payoff of 0 for both partners.

Using these values to develop an ordinal order of preferences for both players the PD is as shown in figure 4. For both players the ordinal order of preferences is $1 > 2 > 3 > 4$. Hence the dominant strategy for both players is strictly not to cooperate. If the other one cooperates, no cooperation is better because it can gain $E_{WLAN,tx}$ and if the other player does not cooperate, no cooperation is still the better option, as one would not have the advantages deriving from cooperation (the chunks the other one is sending) and still has to carry the transmission costs $E_{WLAN,tx}$.

Looking at the total welfare of the interaction (i.e. the combined payoffs of both partners) it becomes obvious that the result from the individual preferences is the worst case, as the total welfare is 0, compared to potentially $2 * (E_{no\ coop} - E_{coop})$ in the case of cooperation. Hence, although an overall beneficial result is possible, individual preferences lead to a non-beneficial state.

From this analysis, it becomes evident that for WMG in particular and NGMN to be successful, there are not only technological, but also economic challenges to meet. Without understanding these economic issues, the work invested in development of different protocols and transmission networks may bring only sub-optimal results. Hence, the longer term aim of our research is the analysis of socio-economic policies to address the issues identified. Consequently, we intend to analyse how normative frameworks and different socio-economic enforcement

⁶The absolute size of this energy difference depends both on the size of the file to download as well as on the number of cooperation partners, that is why figure 3 only uses relative payoffs.

mechanisms affect the cooperative behaviour in WMG and NGMN, as well as other systems, where this kind of cooperation dilemma appear (e.g P2P settings) using design- and run-time models. For us, what makes WMG a particularly interesting object of study and distinguishes the concept from other systems (such as the cooperation problems in P2P settings) are the mobile phone specific resource restrictions (e.g. SIM card space, battery capacity) and features (mobility of the users), as well as its high expected impact in future 4G networks. Thus, current WMG scenarios by the mobile phone industry include cooperative streaming of videos at big sport events, the sharing of news and financial data in banking districts, IPTV, cooperative online gaming as well as cooperative downloads of maps and location information at airports, etc. Finally, due to the comparatively high rates of transmission failures of wireless applications, a further major challenge is answering the question of intention when non-cooperation is observed — was it the result of cheating or transmission failure?

3. A Normative Framework to Explore the Prototyping of WMGs

The essential idea of normative frameworks is a (consistent) collection of rules whose purpose is to describe *a principle of right action binding upon the members of a group and serving to **guide**, **control**, or **regulate** proper and acceptable behaviour [Merriam-Webster dictionary]*. These rules are stated in terms of events, specifically the events that matter for the for the purposes of the normative framework. Thus an important aspect of the definition of a normative framework is how those events are observed. The key concept here is what Searle (1969) terms *conventional generation*, whereby an observable event in the environment is interpreted as a normative event, causing an update to the normative state, as illustrated in Figure 5.

The formalization of this is quite simple, being defined as conditional operations on a sets of terms that represent the normative state. The condition on an operation is typically expressed in terms of a combination of a set of terms and an event, while the operation is an action on the set of terms. One further key concept is the distinction between a environment event and a normative event. This is perhaps best illustrated by an example: that someone dies is an event, but whether the individual was killed or was murdered is a fact that may only be established by a judicial process, which is a normative framework. This notion of attributing a normative label to a environmental event occurs frequently in legal and social frameworks and is known the “counts-as” principle (Jones and Sergot, 1996).

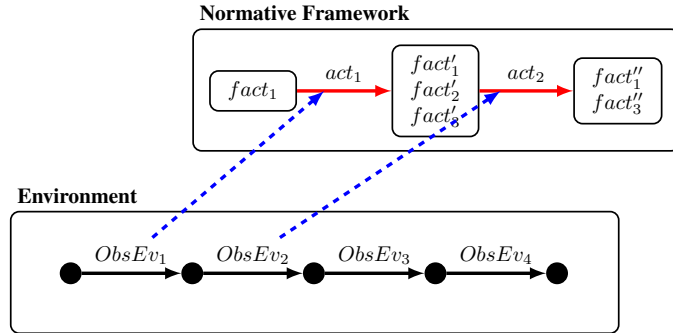


Figure 5: Conventional generation turns observations into normative facts

The essential elements of our normative framework are: (i) events (\mathcal{E}), that bring about changes in state, and (ii) fluents (\mathcal{F}), that characterize the state at a given instant. The function of the framework is to define the interplay between these concepts over time, in order to capture the evolution of a particular institution through the interaction of its participants. We distinguish two kinds of events: normative events (\mathcal{E}_{norm}), that are the events defined by the framework and exogenous (\mathcal{E}_{ex}), that are outside its scope, but whose occurrence triggers normative events in a direct reflection of the “counts-as” principle discussed above. We summarize the elements of the formal model in Table 3.

The semantics of the normative framework is defined over a sequence, called a *trace*, of exogenous events. Starting from the initial state, each \mathcal{E}_{ex} is responsible for a state change, through the initiation and termination of fluents. This is achieved by a three-step process: (i) the transitive closure of \mathcal{G} with respect to a given exogenous event determines all the generated (normative) events, (ii) to this all violations of events not permitted and obligations not fulfilled are added, giving the set of all events whose consequences determine the new state, (iii) the application of C to this set of events identifies all fluents that are initiated and terminated with respect to the current state so giving the next state. For each trace, we can therefore compute a sequence of states that constitutes the model of the normative framework for that trace. This process is realized as a computational model through Answer Set Programming (Baral, 2003; Gelfond and Lifschitz, 1991) and which is the subject of the evaluation process in Section 4.2. Cliffe et al. (2007) shows that the formal model of a normative framework can be translated to an *AnsProlog* program, such that the answer sets of the program correspond to the traces of the framework. A detailed description of the mapping can be found there.

Description	Formalization
Normative system	$\mathcal{N} := \langle \mathcal{E}, \mathcal{F}, \mathcal{C}, \mathcal{G}, \Delta \rangle$
Normative facts (fluents), comprising power, permission, obligations and domain-specific facts	$\mathcal{F} = \mathcal{W} \cup \mathcal{P} \cup \mathcal{O} \cup \mathcal{D}$
Generation relation: maps a state and an event to a set of events	$\mathcal{G} : \mathcal{X} \times \mathcal{E} \rightarrow 2^{\mathcal{E}_{norm}}$
Consequence relation: maps a state and an event to a pair of sets of fluents, where the first are additions and the second are deletions	$\mathcal{C} : \mathcal{X} \times \mathcal{E} \rightarrow 2^{\mathcal{F}} \times 2^{\mathcal{F}}$ where $\mathcal{C}(X, e) = (C^\uparrow(\phi, e), C^\downarrow(\phi, e))$ where (i) $C^\uparrow(\phi, e)$ initiates a fluent (ii) $C^\downarrow(\phi, e)$ terminates a fluent
Events, comprising exogenous, (normative) actions and (normative) violations	$\mathcal{E} = \mathcal{E}_{ex} \cup \mathcal{E}_{norm}$ with $\mathcal{E}_{norm} = \mathcal{E}_{act} \cup \mathcal{E}_{viol}$
The initial set of fluents	Δ
State formula, being the set of positive and negative fluents that comprise the current normative state	$\mathcal{X} = 2^{\mathcal{F} \cup \neg \mathcal{F}}$

Table 3: Formal specification of the normative framework

4. The Design-time Model of the Wireless Mobile Grid Scenario

4.1. The Normative Framework

Now that we have described the wireless grid scenario in some detail from the technological perspective (Section 2) and outlined our approach to modelling normative frameworks (Section 3), we can combine these to present the intuition behind our normative model for wireless grids. The full model specification can be found in Balke et al. (2011).

Implementing a simulation environment is a complex and time-consuming task. Before starting the process it is best to verify that the protocol is indeed suitable for the task at hand. The normative model provides the protocol designer a means to verify, from a theoretical perspective, that the protocol is correct and the methods of enforcing collaboration between the users will indeed benefit all participants in terms of reduced communication costs, battery costs and that none of the participants has an unfair advantage.

The model given here is intentionally limited to focus on the *essential* interactions and the communication costs that arise from those interactions. A more

elaborate model may be desirable for greater realism, but we believe that more details would largely distract and complicate, rather than add to the presentation.

The base-station uses several different frequencies (frequency division multiplexing), allowing many users do download simultaneously. We refer to a frequency division in the model as a channel.

The features of the prototypical scenario are: (i) $1 \times$ base-station: B , with $C = \{c_1, \dots, c_n\}$ channels (ii) $m \times$ users with handsets: $A = \{a_1, \dots, a_m\}$, (iii) $1 \times$ digital good: G divided into, n chunks: $\{g_1, \dots, g_n\}$. We further assume that $n|m$, which is to say the number of chunks, is a multiple of the number of users.

In the design-time model, we focus on the interaction between participants, neglecting the preceding phase of searching for cooperation partners. We identify three phases to the interactions for handset to base-station and handset to handset:

1. Negotiation: assign g_i to a_j s.t. $f : G \rightarrow A$ and

$$f^{-1} : A \rightarrow G \text{ where } f^{-1}(a_i) = \{g_j, \dots\} \text{ and } f(g_j) = a_i$$

2. Downloading: handset a_i receives chunks $f^{-1}(a_i)$ from base-station B
3. Sharing: handset a_1 sends chunks $f^{-1}(a_i)$ to and receives chunks $G \setminus f^{-1}(a_i)$ from other handsets.

While these three phases are distinguishable, they need not be sequential. Of course, the negotiation phase has to precede both the downloading and sending/receiving phases. Sharing is possible as soon as downloading has commenced; thus the two can be interleaved. In the following paragraphs we discuss the normative aspects of each phase in more detail.

Negotiation Phase. In the design-time model, we are not particularly concerned with the technicalities of the negotiation phase — any off-the-shelf protocol could be employed — as long as the post-condition is satisfied: that each chunk is assigned to exactly one handset and that each handset is assigned the same number of chunks. These conditions can readily be relaxed at the cost of a lengthier specification. The pre-defined allocation of chunks for downloading from the base-station is represented in the initial state of the model using fluents to associate handsets and chunks. In addition, each handset has the permission and power to download chunks.

Downloading Phase. This is when each handset downloads its assigned chunks. This process should result in each handset holding $n|m$ distinct chunks. In our model each handset can only physically obtain one chunk at a time from the base station, while each channel can only be used to download a single chunk. A

request to download a chunk is granted whenever there is an available channel and the handset is not currently receiving from the base-station and is not busy sending another chunk. When a chunk is downloaded, the handset and the channel are busy for a fixed amount of time — 4 time steps in this case. From the first instant of the handset interacting with the base-station, it is deemed to have downloaded the chunk, so parts can be shared. As soon as a channel and a handset are engaged, the framework (i) removes the power from the handset and from the channel to engage in any other interactions, (ii) stops the handset from needing the chunk and (iii) cancels the permission to download the chunk again later on. We model the passage of time by generating a transition for each exogenous event. But, in case no handset is interacting with the normative framework, we use a clock event to achieve the same effect. Likewise, we use a transition event to count down the elapsed time channel/handset interaction. Thus, when an interaction comes to an end, power is restored to the handset to download chunks via the channel and for the handset to receive more chunks.

Sharing Phase. In this phase each handset sends chunks to, or receives chunks from other handsets. At the end of the process, each handset has a complete set of the chunks, that is the entire digital good. The idea is similar to that for with downloading chunks from the base-station, only that we utilize a simple mechanism to encourage handsets to share their chunks with others instead of downloading them. Just as for downloading we must track handset activity, whether they busy and whether they are sending or receiving. A handset sharing a chunk can possibly trigger two normative events, indicating sending or receiving, depending on the availability of the handsets and whether they have possession of the chunk or not. The duration for sending or receiving a chunk is set to 2 time steps. The time steps chosen for downloading and sending/receiving have been chosen arbitrarily to demonstrate the difference between the two communication mechanisms. They were kept relatively small to reduce the length of the traces while still allowing for possible simultaneous receiving and downloading. When a chunk is received from another handset, the handset loses permission to receive another chunk until it has sent a chunk to another handset. Continuous receiving without sending results in a violation of the protocol. The penalty applied is that the violating handset permanently loses the power to receive, which means that for all intents and purposes it has been expelled from the collaboration group.

Running the Model. When we translate the model into *AnsProlog* and run the program, we obtain all the possible traces over a specified number of time instances. A successful trace is defined as one in which all handsets have all chunks

and are no longer engaged. Figure 6 shows the graphical representation of a trace for a scenario with two handsets (bob and alice), four chunks (x1, x2, x3 and x4) and a base-station with two channels (c1,c2). The little circles indicate the time steps. If the circle is light grey this means the handset is busy receiving from the base-station, while darker grey indicates it is busy sending or receiving to/from another handset. Circles half light, half dark indicate the handset is both receiving from the base-station and another handset. The arrows indicate which chunk goes to which handset. The left-hand side labels indicate the exogenous event and the current distribution of chunks.

4.2. Evaluation

With the normative design-time model, we can examine the traces for expected and unexpected behaviour. Furthermore we can obtain an estimate for communication cost (the energy consumption) in each trace simply by counting the number and type of communication events for each handset, then using the corresponding information from Tables 1 and 2.

At this stage the model is used as a design-time tool and generates all possible traces. As mentioned earlier, we are only interested in those traces that lead to success. Even so constrained, the model still generates a large number of valid traces. But on closer analysis, it becomes clear that the occurrence of most of these traces is a function of the relative intelligence and (bounded) rationality of the handset. For example, repeatedly trying the download a chunk when the handset is busy. The model purposely does not attempt to characterize handset behaviour—we believe that is responsibility of the handset designer—because our objective is the design of the space in which the handsets interact. However, these unnatural traces can easily be filtered out by additional constraints on the *Ans-Prolog* specification, so that it retains only those traces in which an exogenous event leads to an normative event, reducing the number of traces significantly.

When verifying the model we distinguish between two types of traces: those that do not contain any violations and those that do. This can be easily accomplished by adding further filters on the traces.

Our model has two types of violations, downloading unassigned chunks and receiving without sharing. In our current model no sanction is imposed on the former type of violation, as handsets are already penalised with a lower data rate. In the current model it is possible for a handset to take advantage of the protocol by violating it in order to obtain the last chunk after which the sanction does not have any effect. This is a typical situation in a game theoretic context, resulting

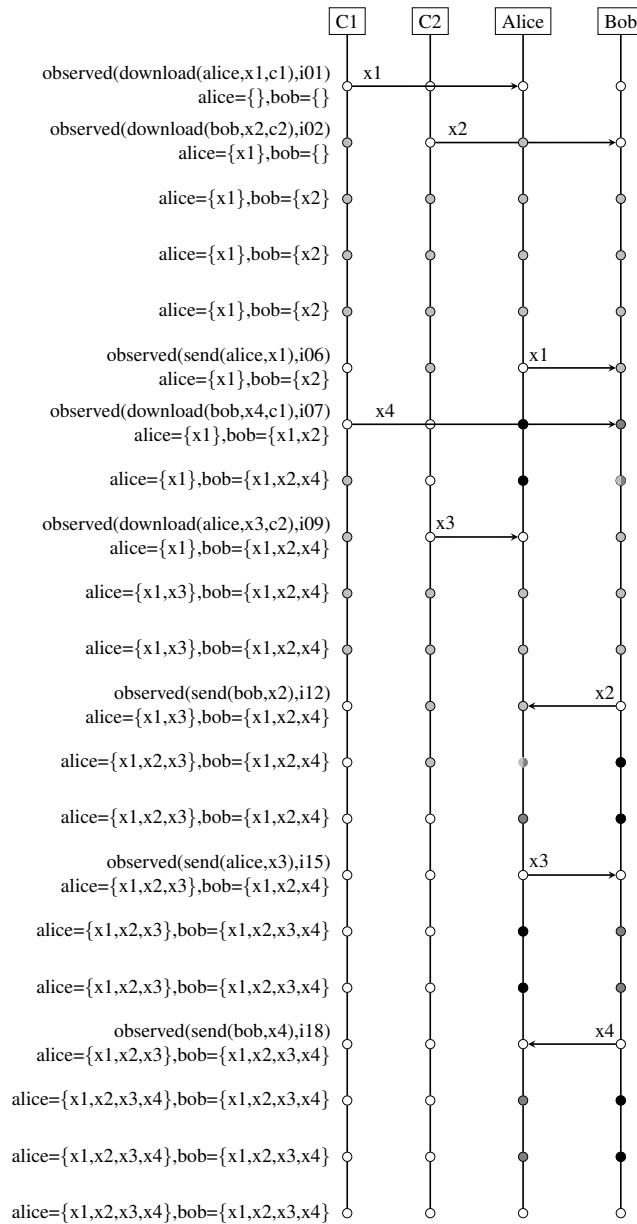


Figure 6: One trace of the interaction between alice, bob and the channels of base-station

in the system reverting back to non-collaborative behaviour. To overcome this, a reputation mechanism could be used to penalise the handset in future interactions.

The current model takes a rather harsh position on sanctioning, in that the violating handset is expelled — the power to receive chunks from another handset is rescinded. In fact, this is both harsh and counter-productive, because given the initial state, the chunk assignment is not 1-resilient — meaning the distribution cannot be achieved in the case of the expulsion of 1 handset, except in the special case where the expulsion occurs when the other handsets no longer require any chunks from this one. Full 1-resilient assignment can be achieved with two chunks for each of three handsets, in which each chunk is assigned to two handsets and of course, n -resilience can be achieved by each handset downloading all the chunks from the base-station. In terms of the effect on the group goal, the ejection scenario is equivalent to one of the handsets leaving the ad-hoc network. In either case (expulsion or departure), for an a-priori solution there is a trade-off to be explored in delivering i -resilience, based on the estimated number of handset failures and the additional cost of replicated base-station downloads. Alternatively, some handsets may engage autonomously in additional base-station downloads for the sake of the group goal. A more practical sanction may be to lock the offending handset out of the sharing process for a number of time steps, but as discussed above, the appropriateness of this depends on the impact on the group goal.

The model also assumes alternating sending and receiving. In reality this might not always be the case. Handsets should be allowed to take advantage of chunks being sent even when the same number of chunks have not yet been shared. It might be more realistic to evaluate a handset's willingness to collaborate over a larger time period.

5. Simulating the Wireless Mobile Grid using an MAS Simulation

5.1. NGMN as Multi-Agent-Systems

Having presented the design time model, we now follow it with a runtime model, realized as simulation. The purpose of the runtime model is to allow us to explore the scenario from the perspective of the individual decision makers (handsets) and as a consequence collect data about the overall system behaviour. For the runtime model simulation we chose a Multi-Agent-System (MAS) simulation, using the well-established notion of agents as a metaphor to represent human users with their mobile phones in the simulation. Citing Ferber (1999) an agent is “a physical or virtual entity that (i) is capable of acting in an environment, (ii) can communicate directly with other agents, (iii) is driven by a set of tendencies (in the form of individual objector of a satisfaction/survival function which it tries to optimise), (iv) possesses resources of its own, (v) is capable of

perceiving its environment (but to a limited extent, (vi) has only a partial representation of this environment (and perhaps none), (vii) possesses skills and can offer services, (viii) may be able to reproduce itself, (ix) has behaviour tending towards satisfying its objectives, taking account of the resources and skills available to it and depending on its perception, its representations and the communications it receives.”

Keeping this definition in mind, in a NGMN a user together with his mobile phone (as one entity) can be seen as an agent. Thus, the users with their phones act in the environment and interact with other users, each being driven by their own objectives (file download and energy saving) as well as being constrained by resource limitations (battery capacity and bandwidth). The actions the users perform are, on the one hand based on their perception of their environment, (e.g. of the other users and their actions) and on the other on their resources (e.g. files they have and need) as well as utility considerations (battery costs for different actions).

A MAS is a system composed of multiple interacting agents that all make their own decisions based on their own goals, utilities and resources. Thus, in contrast to other forms of simulation the overall system behaviour is not encoded by the simulation designer in advance, but the macro behaviour results from the individual decision and actions of the agents (i.e. a micro-level driven macro result). Hence, using MAS simulation for simulating the actions and interactions of autonomous agents one gets a chance to assess the macro effects on the system as a whole that result from micro-behaviour of the individual agents, something that is particularly suitable for the problems discussed earlier.

What happens in the course of the simulation itself is that at each step, agents receive an observation from the environment, change their internal state and take an action — which may even be to do nothing — that is finally executed. The internal state of an agent possibly encodes its history of actions and observations, its beliefs about the state of the environment, as well as its own preferences. The internal state evolves by integrating observations from its environment. The agent’s decision function reflects its behaviour or policy and determines which action it will take in the next step.

As indicated earlier, one of the most important features of MAS simulation is how the agent actions lead to observable system level macro behaviour. It is for this reason that, in order to explain the simulation in more detail, that we first look at the actions an agent can perform (see section 5.2), in which we observe that we explicitly allow for cheating by the agents, i.e. leave it as an action-option to them. The implementation of these actions as well as the agents and the environment was

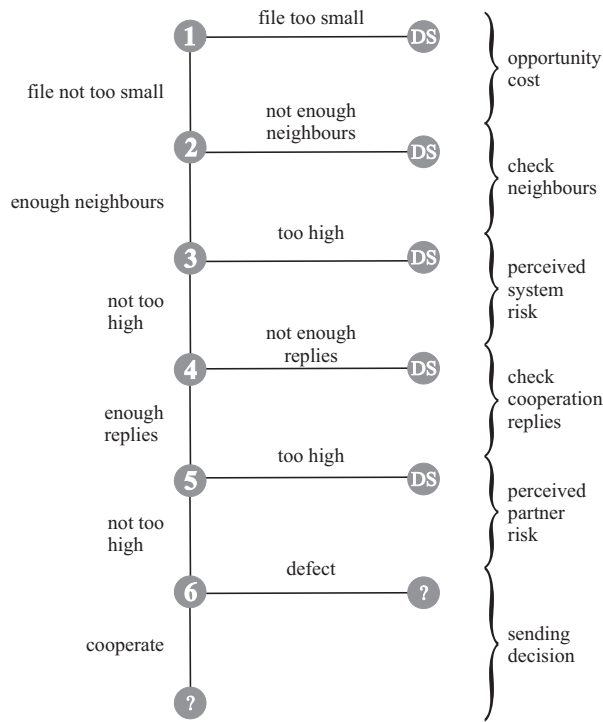


Figure 7: Download Considerations by the Participants

realized using the Jason Simulation Platform (Bordini et al., 2007).

5.2. The Agent Decision Process

This section serves to explain the decision process of the agents in the simulation. Starting from Figure 7 that shows the agent's a decision tree, we discuss how an agent decides in favour of a particular action. The basis for decision-making are simple utility considerations, but how it arrives at different utility values will be now be explained together with an illustration of the nodes in the decision tree.

For ease of reference, we have numbered the decision nodes in Figure 7 from 1 to 6, and will refer to these when discussing the respective nodes. When an agent wants to download a file, it sequentially goes through the nodes 1 to 6 and every time decides whether it should download the files itself (DS) or whether to go on to the next decision node.

Node: 1 Wanting to download a file, the agent starts at node 1 and considers the

opportunity costs of the download. These are the battery costs accumulating for searching for and negotiating with potential cooperation partners. If these costs are higher than the gain from a cooperative download (i.e. higher than the saved energy if the cooperation is successful), the agent will download the file itself. Since the energy saving from a cooperative download depends on file size, we compare this information (and the resulting minimum energy gains in case of a successful cooperation) against the minimum energy needed for sending 1 cooperation request and 1 negotiation message. Hence, if the file is large enough for the cooperative download energy saving to outweigh the opportunity costs, the agent will move to node 2 (otherwise it will simply download the file itself).

Node: 2 At node 2, the next consideration for the agent is the perceived system risk. In the simulation, this risk has a value between 0 and 1 (with 0 being no risk and 1 being the highest risk possible). The perceived system risk is that which the agent has with regard to the environment in which it is interacting. To illustrate: if a system is highly monitored, rules to counter defection are in place and furthermore, there is a high chance of seeing cooperation partners again, the perceived system risk will be low; in contrast in an open distributed system without any enforcement possibilities (such as in the case of the wireless grids if no cooperation ensuring measures are taken) it will approach 1. The higher the perceived system risk (i.e. the closer the value is to 1), the more likely the agent will download the file itself. In the opposite case, it will go on to decision node 3 where the number of neighbours is determined⁷.

Node: 3 This step is straight forward and is an environmental enquiry rather than an actual decision. If this number is too low (e.g. if it cannot find many other agents in its vicinity) then sending a cooperation request has little point, since it only costs battery and will not result in enough responses, so the agent downloads the file itself. If the number of neighbours is high enough, it does the opposite, i.e. sends out a cooperation request⁸

⁷The decisions in node 2 and 3 need not necessarily be made in this order but could be made simultaneously, for reasons of simplicity we however assume sequential ordering.

⁸In the simulation implementation a further case is distinguished: the agent receives a cooperation request from another agent in its vicinity. If this request matches its interest, i.e. the request

to its neighbours via an IEEE802.11 WLAN broadcast and waits for responses.

Node: 4 Decision node 4 is similar to node 3, only this time the agent checks for the number of responses to its cooperation message rather than the number of neighbours.

Node: 5 Having received enough responses in decision step 5, the agent checks who has responded and decides whether it wants to work with the other agents. Thus, in reality an agent might get a response from an agent that has cheated on it before or which for other reasons does not seem trustworthy in its eyes and therefore could decide not to cooperate with this agent.

Node: 6 Having agreed to join a cooperation group, finally, in decision step 6, the agent has two tasks. First of all it has to decide whether to download its promised share for the cooperation group from the base-station and more importantly, having downloaded its share, decide whether to cooperate (i.e. send its share to its cooperation partners) or to defect (i.e. not send its share). This decision is based on the utilities from the Prisoner's Dilemma described earlier as well as possible follow-up costs. Thus, if it defects but is likely to be levied a fine under these circumstances, the original utilities from the PD change and it might be advantageous to cooperate instead. In contrast, with no follow-up costs the original PD will be relevant and defection will always be the better option.

Having made its decision, the last step (which however is not a decision and so does not appear in Figure 7) is to wait and see whether the cooperation partners have sent their promised shares. If shares are missing, the agent has to go through the decision process again and can decide whether to download the missing shares itself or whether to try to find new cooperation partners.

5.3. Implementation

5.3.1. The Agent Decision Process

Having explained the basic utility-based download considerations of the agents, in this section, we describe their implementation as well as other important parts

is for the same file the agent wants to download, it does not send a request itself, but only answers the other request. For reasons of simplicity this case is omitted from the decision tree and it is assumed that it has not received any matching requests.

of the simulation.

Jason uses AgentSpeak (Rao, 1996) for describing agent behaviour, thus following the Beliefs – Desires – Intentions (BDI) model of agency (Rao and Georgeff, 1995). BDI architectures originated in the work of the Rational Agency project at Stanford Research Institute in the mid-1980s. The main idea of BDI is that computer programs (i.e. the agents in the simulation) are viewed as if they have a “mental state”. Thus, when programming the agents, computational analogues of beliefs, desires and intentions are used, enabling a form of reasoning by agents about their goals and the different options to achieve them.

Beliefs are information the agent has about the world. This information could be out of date or inaccurate, of course, however they represent the agent’s view of the world. Desires are all the possible states of affairs that the agent might like to accomplish. Having a desire, however, does not imply that an agent acts upon it: they only potentially influence an agent’s actions. Finally, intentions are the states of affairs that the agent has decided to work towards. Intentions may be goals that are delegated to the agent, or may result from considering options: we think of an agent looking at its options and choosing between them. Options that are selected in this way become intentions. Therefore, one can imagine an agent starting with some delegated goal, and then considering the possible options that are compatible with this delegated goal; the options that it chooses are then intentions, to which the agent is committed (Bordini et al., 2007).

Jason makes use of the BDI concept by repeatedly executing the following control loop: (i) the individual agents look at the world, perceive their environment and other agents, and update their individual beliefs on this basis (it is important to note that not all agents perceive the same, but have individual percepts that can be different for each agent); (ii) as a result they deliberate to decide which individual intention to achieve; (iii) and use means-ends reasoning to find a plan (a step of actions) to achieve this intention; (iv) in the last step the agents then execute the plan in order to fulfill the intention.

In our simulation all environmental related aspects as well as the mathematical calculation of utilities by the individual agents are programmed in Java, whereas the agents’ reasoning about goals and actions is implemented in AgentSpeak. In addition, the Java component of Jason was used for logging the states of the MAS.

5.3.2. *Time*

An important aspect of any system designed to model action and change is how it deals with the problem of representing time. As is pointed out (Allen, 1991) a number of methods have been found to be useful in different areas of

artificial intelligence, including representations based on explicit dating, intervals, and temporal logics (Cliffe, 2007). For the purposes of this paper we assume a model of time consisting of a set of totally ordered time instants for the agents' actions, such that at each point exactly one "real-world" event may occur. In this respect we do not make any assumptions about durations of real time which may take place between these instants. With regard to the download times and the resulting energy consumption, we use the file size and the data-rates (and power values) given in Tables 1 and 2 to determine the respective values regardless of the model of time consistency.

5.3.3. *The Notion of Location*

One last implementation aspect that will be discussed briefly, is that of location. This is also important, as the movement of agents in the model and their respective location in the system and in relation to other agents determines potential cooperation options.

For cooperation to take place, an agent must be in the proximity of other agents. An accurate Cartesian model of location is not actually necessary, so we represent an agent's location as a number between 0 and 1 and proximity is determined by a search interval $[location - x, location + x]$, which the agent uses to find cooperation partners.

5.3.4. *Simulation Variables and Results*

Having discussed the implementation aspects of the MAS simulation, in this section we now describe the simulation experiments and the results collected.

For the simulation experiments we used a very similar scenario as in the design time case, but without an enforcement mechanism. Table 4(a) gives an overview of the parameters that can be used to control the simulation⁹ at the start of the experiment.

The simulation parameters comprise, the number of the agents in the simulation, the duration of the simulation (i.e. the number of interactions for which the simulation is running), the size of the partner search interval that is used for proximity determination as well as the location-based information, i.e. the mobility model and the mobility levels of the agents. The mobility level is a number

⁹The variables that result from agent decision or are agent specific such as perceived risks, etc. are not part of the simulation parameters as they cannot be altered directly. The perceived system risk was artificially influenced externally for one set of experiments, which is why it is listed in parentheses in Table 4(a).

Table 4: Simulation Specifications

(a) Simulation Variables		(b) Simulation Parameters	
Name	Range/Type	Name	Simulation Parameters
Agents	[2, ∞]	Agents	2000
Interactions	[1, ∞]	Interactions	50
Mobility Level	[0, 1]	Mobility Level	uniform distribution
Movement Model	Random Walk, Levy Flight,...	Movement Model	Random Walk
Partner Search Interval	[0, 1]	Partner Search Interval	0.1
(Perceived System Risk)	[0,1]	(Perceived System Risk)	0, 1
Enforcement	[true, false]	Enforcement	[true, false]

between 0 and 1 that is assigned individually to each agent. It controls how far an agent can move. The closer the value is to 0, the less distance an agent can move, whereas a level of 1 indicates moving long distances. The idea is that elderly people for example will move around the same place whereas business people might move further on a regular basis (e.g. to meetings) for example.

For the simulation run we used a uniform distribution of the mobility level between all agents. Further parameters of the simulation runs can be seen in Table 4(b).

For the movement pattern, we started with a random walk. Thus, taking mobility level into account, a random number between 0 and 1 is drawn in order to determine the new location. The interaction count was set at 50, while the size of the agent population was set at 2000. Finally, the partner search interval was set to 0.1 and the perceived system risk was to be determined by all agents themselves or set to 0 and 1 as indirect fixed values.

In order to test the effects of the agent decision on the cooperation and as a consequence on battery consumption, we have executed fractional factorial experiments. A factorial experiment is an experiment whose design consists of two or more factors, each with discrete possible values or “levels”, and whose experimental units take on all possible combinations of these levels across all the factors. Our experiments consisted of running a number of simulations (50) for the parameter combinations in Table 4(b). The logfiles of these runs were then analysed using Matlab. The most interesting results of these experiments are shown in Figure 8.

The figures show the boxplots of the average cumulative energy consumption by all agents in different simulation experiments. The results of these simulation experiments were logged and the resulting datasets analyzed with Matlab. In the simulation experiments resulting in datasets 1 and 2 the perceived system risk was set to 0, whereas in the simulation corresponding to dataset 3 it was set to 1. As a result of this, in the simulations for datasets 3, the system perception by the agents

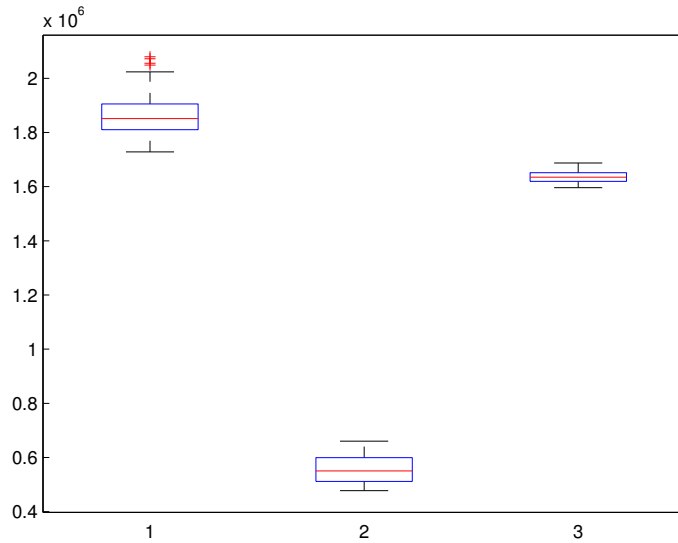


Figure 8: Battery Costs in the Simulation

is classified as too risky (as no enforcement is in place at all) and as a consequence the agents already at node 2 in the decision tree, choose to download everything themselves. What distinguishes the experiments responsible for datasets 1 and 2 is the usage of enforcement. Thus, in the first simulation experiment no enforcement was being used whereas in the second set of experiments we applied enforcement by altering economic considerations by agents, as also proposed by León et al. (2009). Thus, we introduced a very high sanctioning fee, that made cheating very expensive for agents and therefore lead to them considering cooperation instead of cheating.

From the results, one can clearly see that in case 1, i.e. when no enforcement is in place and the system risk is perceived to be low, the energy consumption costs are highest. They are even higher than case 3, where due to the perceived high system risk no cooperation is considered by the agents. The reason for this effect is that the agents as a result of the low perceived system risk in case 1, enter negotiation, for which energy is needed to agree the negotiation group as well as on the file to download for example. However when the actual sharing is supposed to take place, as a result of the utility considerations and no enforcement mech-

anism being in place, the agents decide to cheat, resulting in them not receiving any chunks from cooperation partners and having eventually to download everything themselves as a result. Thus, no energy saving is possible from the short link cooperation and furthermore they have to face additional battery costs for the negotiation.

One further interesting result appears when comparing case 2 with the other results. Case 2 includes a basic enforcement concept in form of a fee (sanction) being applied if agents are caught cheating. It therefore shows the potential gains (in terms of battery) possible in a wireless grid, if enforcement is possible and successful. Thus in case 2 the energy consumed is approximately 3.37 times less than in case 1 and approximately 3.019 times less than in case 3, implying that given the specific setting of the simulation experiments, a 70% energy saving (with regard to case 1) could be achieved.

5.4. Evaluation

With the help of the MAS simulation, we are able to observe the individual agent decisions as well as the impact of these decisions on the system as a whole (i.e. the macro level). As expected from the economic analysis in Section 2.2 the results of the simulation show that without any enforcement mechanism in place the theoretical advantages of wireless grids are at stake and might even result in the opposite, i.e. higher energy consumption due to additional negotiation messages. These results were stable throughout the factorial experiments and are good reference point for exploring enforcement in wireless grids further.

The results indicate the importance and the potential impact of successful enforcement in wireless grids, and any NGMN of this kind in the future, and can be extended to give valuable empirical evidence of the potential impact of enforcement mechanisms at an early prototyping stage. This is particularly important in the early prototyping stage, when wireless grids are not being actually deployed at large and the simulation may uncover possible weaknesses and strengths before costly tests with real people and hardware are conducted.

6. Conclusions and Future Work

In this paper, we have presented a normative framework and a MAS simulation as complementary mechanisms to help understand and to model the economic challenges that might arise in the context of a specific NGMN scenario, namely wireless grids.

Using the design time model of a normative frameworks, we were able to quantify communication costs for the particular case of a 3G structured network and an Ethernet ad-hoc network as well examine the traces for expected and unexpected behaviour. Thus, with the help of the design time model, system properties can be checked as well as validated. Whereas the design time model focuses on the system and its properties as a whole and reasons about the validity of system states, the runtime model complements this approach. Hence, the MAS simulation gives insights into the individual behaviour and decision-making of the agents at the runtime of the simulation and thus allows for a micro-macro view. “Micro”, meaning the individual decision making by the agents, that however culminates in a “macro” result, i.e. a global system-wide emergent behaviour.

Both, the runtime as well as the design time model are currently necessarily simplified and demand expansion, as pointed out in Sections 4.2 and 5.4. For example, in the MAS simulation, points that need development include the size of the agent population and the movement model. Further aspects of interest to be included in the model are error rates on the different communication paths as well as more elaborate models for the negotiation (e.g. some agents having some parts of the files already).

The modelling of the wireless grid scenario gave us also a good insight into our formal model. The model does not allow us to expel an agent completely from a normative framework, as all the observed events are automatically empowered. While this can be partially remedied by removing the empowerment of consequent normative events, as we have done in the sharing phase, it raises interesting issues on how membership of a normative framework should be handled.

One important aspect that needs to be investigated more deeply in both models are the enforcement mechanisms. As presented in this paper, the design time model makes use of an enforcement mechanism that assumes the possibility to observe all actions by the agents and that the normative framework has the power to carry out the sanction of not allowing downloads, if the handset has not shared chunks. The runtime model in contrast does employ enforcement mechanism by threatening agents with such high sanctioning costs, that they will most likely aim for the cooperation option. It goes not into detail on how the enforcement works in detail and does not compare different enforcement approaches, but demonstrates the necessity for a more in depth analysis of possible enforcement mechanisms. Also, in their specific implementation, both models have given valuable insights into the cooperation problem in wireless grids and extending them with realistic (i.e. more subtle, and less draconian) enforcement mechanisms. One first step in this direction has been taken, in form of a classification of possible different

enforcement options, was presented at a Dagstuhl seminar Balke (2009). With the help of the two complementary modelling approaches, we will implement the different enforcement mechanisms in the models and compare and evaluate the success and utility of them.

One particular focus in this comparison/evaluation will be the effects the mechanisms have on the different stakeholders of the system. The point is that if such wireless grids are realized, it is not only users that will be affected, but also the infrastructure providers and mobile phone manufacturers. As these groups may have different goals with regard to wireless grids, their perception of the success or utility of an enforcement mechanism might vary. The need to balance the different interests of these stakeholders therefore displays interesting aspects for future research.

References

- Allen, J., 1991. Time and time again: The many ways to represent time. *International Journal of Intelligent Systems* 6 (4), 341–356.
- Alliance, N., 2006. Next generation mobile networks – beyond hsdpa & evdo. Tech. rep., http://www.ngmn.org/uploads/media/Next_Generation_Mobile_Networks_Beyond_HSPA___EVDO.pdf.
- Axelrod, R., 1981. The emergence of cooperation among egoists. *The American Political Science Review* 75 (2), 306–318.
- Balke, T., 2009. A taxonomy for ensuring institutional compliance in utility computing. In: Boella, G., Noriega, P., Pigozzi, G., Verhagen, H. (Eds.), *Normative Multi-Agent Systems*. No. 09121 in Dagstuhl Seminar Proceedings. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Germany.
- Balke, T., De Vos, M., Padget, J., Fitzek, F., 2011. Using a normative framework to explore the prototyping of wireless grids. In: *COIN 2010*. Vol. 6541 of LNAI. Springer, pp. 97–115.
- Baral, C., 2003. *Knowledge Representation, Reasoning and Declarative Problem Solving*. Cambridge Press.
- Bordini, R. H., Hübner, J. F., Wooldridge, M., 2007. *Programming Multi-Agent Systems in AgentSpeak using Jason*. Wiley Series in Agent Technology. John Wiley & Sons.

- Cliffe, O., 2007. Specifying and analysing institutions in multi-agent systems using answer set programming. Ph.D. thesis, University of Bath.
- Cliffe, O., De Vos, M., Padget, J., 2007. Specifying and reasoning about multiple institutions. In: Noriega, P., Vázquez-Salceda, J., Boella, G., Boissier, O., Dignum, V., Fornara, N., Matson, E. (Eds.), COIN 2006. Vol. 4386 of Lecture Notes in Computer Science. Springer Berlin / Heidelberg, pp. 67–85.
- Ferber, J., 1999. Multi-Agent System: An Introduction to Distributed Artificial Intelligence. Addison Wesley Longman.
- Fitzek, F. H. P., Katz, M. D., 2007. Cellular controlled peer to peer communications: Overview and potentials. In: Fitzek, F. H. P., Katz, M. D. (Eds.), Cognitive Wireless Networks. Springer, pp. 31–59.
- Gelfond, M., Lifschitz, V., 1991. Classical negation in logic programs and disjunctive databases. *New Generation Computing* 9 (3-4), 365–386.
- Jones, A. J., Sergot, M., 1996. A Formal Characterisation of Institutionalised Power. *ACM Computing Surveys* 28 (4es), 121, read 28/11/2004 http://www-lp.doc.ic.ac.uk/_lp/Sergot/InstitPower.ps.gz.
- Katz, M. D., Fitzek, F. H. P., 2006. Cooperation in 4g networks - cooperation in a heterogenous wireless world. In: Fitzek, F. H. P., Katz, M. D. (Eds.), Cooperation in Wireless Networks: Principles and Applications. Springer, pp. 463–496.
- León, X., Trinh, T. A., Navarro, L., November 2009. Using economic regulation to prevent resource congestion in large-scale shared infrastructures. *Future Generation Computer Systems* 26 (4), 599–607, elsevier Science Publishers B. V.
- Ostrom, E., June 1999. Coping with tragedies of the commons. *Annual Review of Political Science* 2, 493–535.
- Perrucci, G. P., Fitzek, F. H., Petersen, M. V., 2009. Energy saving aspects for mobile device exploiting heterogeneous wireless networks. In: Heterogeneous Wireless Access Networks. Springer US.
- Rao, A. S., 1996. Agentspeak(L): Bdi agents speak out in a logical computable language. In: MAAMAW '96: Proceedings of the 7th European workshop on Modelling autonomous agents in a multi-agent world: agents breaking away. Springer-Verlag New York, Inc., pp. 42–55.

- Rao, A. S., Georgeff, M. P., 1995. Bdi-agents: from theory to practice. In: Proceedings of the First International Conference on Multiagent Systems.
- Searle, J. R., 1969. Speech acts: An essay in the philosophy of language. Cambridge University Press.
- Wrona, K., Mähönen, P., 2006. Stability and security in wireless cooperative networks. In: Cooperation in Wireless Networks: Principles and Applications – Real Egoistic Behavior is to Cooperate! Springer Netherlands, pp. 313–363.