Axiomatisation of Socio-Economic Principles for Self-Organising Institutions: Concepts, Experiments and Challenges

JEREMY PITT, Imperial College London
JULIA SCHAUMEIER, Imperial College London
ALEXANDER ARTIKIS, National Centre for Scientific Research 'Demokritos'

We address the problem of engineering self-organising electronic institutions for resource allocation in open, embedded and resource-constrained systems. In such systems, there is decentralised control, competition for resources and an expectation of both intentional and unintentional errors. The 'optimal' distribution of resources is then less important than the endurance of the distribution mechanism. Under these circumstances, we propose to model resource allocation as a common-pool resource management problem, and develop a formal characterisation of Elinor Ostrom's socio-economic principles for self-governing institutions. This paper applies a method for sociologically-inspired computing to give a complete axiomatisation of six of Ostrom's eight principles in the Event Calculus. A testbed is implemented for experimenting with the axiomatisation. The experimental results show that these principles support enduring institutions, in terms of longevity and membership, and also provide insight into calibrating the transaction and running costs associated with implementing the principles against the behavioural profile of the institutional membership. We conclude that it is possible to express Ostrom's principles in logical form and that they are necessary and sufficient conditions for enduring self-organising electronic institutions to manage sustainable common-pool resources.

Categories and Subject Descriptors: I.2.11 [**Distributed Artificial Intelligence**]: Multiagent Systems

General Terms: Algorithms, Performance

Additional Key Words and Phrases: Norm-Governed Multiagent Systems, Self-Organising Systems, Institutions, Resource Allocation, Common-Pool Resource

ACM Reference Format:
Pitt, J., Schaumeier, J. and Artikis, A. 0000. Self-Organising Electronic Institutions for Common-Pool Resource Management. ACM Trans. Autonom. Adapt. Syst. 0, 0, Article 0 (0000), 39 pages.
DOI = 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

1. INTRODUCTION

Applications in which autonomous and heterogeneous agents form opportunistic alliances, requiring them to share collective resources in order to achieve individual objectives, are increasingly common. Examples include vehicular networks [Raya and Hubaux 2007], service-oriented systems such as cloud computing [Ardagna et al. 2011], and demand-side infrastructure management for water [Boulet et al. 2009], energy [Strbac 2008], and so on. These examples are all open, distributed and resource-constrained. This constraint on resources means we are unable to 'privatise' the sys-

---

Author's addresses: J. Pitt, Department of Electrical and Electronic Engineering and Institute for Security Science and Technology, Imperial College London, SW7 2BT, UK; J. Schaumeier, Department of Electrical and Electronic Engineering, Imperial College London, SW7 2BT, UK; A. Artikis, National Centre for Scientific Research 'Demokritos', Athens 15310, Greece.

DOI 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

tem, otherwise it would no longer be open, nor can we ‘centralise’ the system, otherwise it would no longer be distributed.

Instead, we address the issue of resource constraint from the perspective of self-governing institutions for common-pool resource (CPR) management [Ostrom 1990]. By this definition, an institution embodies the rules which specify the conditions concerning the provision and appropriation of resources. These rules should be mutable by other rules, and can so be adapted to suit the environment in which the system is embedded. The environment itself is changed by actions of the agents and might also be changed by external forces.

The institution then has to satisfy three performance criteria. Firstly, the coordination mechanisms and conventions should encourage compliance pervasion, defined as behaviour in accordance with the rules or norms, amongst members of the institution. Secondly, the selection, modification and adaptation of the rules must be based on collective decision-making and tolerance of unintentional errors, and should not only suit the environment but also result in a ‘satisfactory’ outcome for each of the members, given an appropriate metric metric for measuring ‘satisfactory’, for example ‘fairness’ [Lan et al. 2010; de Jong and Tuyls 2011], or ‘efficiency’ (e.g. ‘utilitarian social welfare’ [Chevaleyre et al. 2007]). Thirdly, any ‘satisfactory’ distribution of resources has to be sustainable; in other words, the rules also have to ensure that the institution itself is *enduring*, in such a way that the goals of a sustainable resource and an enduring institution are co-dependent.

In this paper, we apply a method for sociologically-inspired computing to examine three propositions: firstly (**p1**), that open, embedded and resource-constrained systems can be considered from the perspective of enduring institutions for management of common-pool resources (CPRs); secondly (**p2**), that socio-economic principles for enduring institutions can be considered from the perspective of norm-governed systems, and can be axiomatised using action languages used in Artificial Intelligence for reasoning about action, agency and norms; and thirdly (**p3**), that such an axiomatisation can be used as a formal specification to implement a testbed to conduct experiments to test whether these principles are necessary and sufficient conditions for *enduring* electronic institutions.

The demonstration of these propositions is organised as follows. In Section 2 the problem of resource allocation in embedded systems with exogenous and endogenous resources is introduced. Section 3 reviews the background work against which these propositions are to be tested, namely the work on CPR management of Ostrom [1990], the concept of institutionalised power [Jones and Sergot 1996], and the framework for dynamic specification of norm-governed systems [Artikis 2011]. Section 4 concerns proposition **p1** and shows how Ostrom’s tripartite analysis of institutions for CPR management can be represented as a dynamic specification. Section 5 addresses proposition **p2** by giving an indicative axiomatisation of six of Ostrom’s eight principles in the Event Calculus [Kowalski and Sergot 1986]. In Section 6, a testbed is described for experimenting with the axiomatisation, experimental results are presented which demonstrate proposition **p3**. These results are then compared by their ‘degree of efficiency’ and used to provide insight into calibrating the transaction and running costs associated with implementing the principles against the behavioural profile of the institutional membership. We discuss related work and further research challenges in Section 7, including possible applications in Cloud Computing and SmartGrids, before concluding in Section 8 with the observations that by applying the method for sociologically-inspired computing it is possible to express Ostrom’s principles in logical form, and that they are indeed necessary and sufficient conditions for enduring self-organising electronic institutions to manage sustainable common-pool resources.

2. INFORMAL PROBLEM SPECIFICATION

In this section, we informally specify the problem of resource allocation in open, embedded and resource constrained systems. We distinguish between systems with exogenous resources, where the resources are derived from an environment over which the system components (which we will henceforth refer to as *agents*) have no control, and systems with endogenous resources, where the agents themselves must supply the resources for the system to operate.

2.1. Resource Allocation in Open Systems

Open embedded systems consist of heterogeneous components of unknown provenance that are coordinating their behaviour in the context of an environment which may be perturbed by outside events. Such systems arise in a class of wireless network, for example mobile ad hoc, opportunistic, sensor and vehicular networks; in service-oriented systems like virtual organisations and cloud computing applications; and increasingly in demand-side infrastructure management, for water, energy, and so on.

All these applications share a number of features. Primarily, decision-making is too fast, frequent and complicated for manual operator intervention. Therefore the system has to be able to operate autonomously. Being open, there is no central controller, no common goal and no common knowledge: therefore collective decisions must be made in the face of both uncertainty and possibly conflicting opinions and requirements. Openness also implies the system must operate in expectation of error, non-compliance to the specification and other sub-ideal behaviour, including both intentional and unintentional violations. Moreover the system components cannot expect any level of cooperation, i.e. that appropriate action will be taken to recover from errors or sub-ideal states. Finally, the system is resource-constrained and the agents are required to share and appropriate resources in order to satisfy individual goals.

We will assume that the sharing and appropriation of resources occurs in discrete time intervals called *timeslices*. We can then distinguish between three different types of system: systems with exogenous resources, systems with endogenous resources, and hybrid systems which have both.

In exogenous systems, the resources are derived from an environment over which the agents have no control. The sequence of operation in a timeslice is that agents first request resources for the current timeslice, the system allocates resources to the agents, and finally the agents appropriate the resources. A typical example is a water distribution system. In endogenous systems, the agents themselves must supply the resources for the system to operate. The sequence of operation is that agents first contribute resources to the common pool, and then proceed with the request–allocate–appropriate sequence as before. Typical examples include MANET and sensor networks. In hybrid systems, there are both endogenous and exogenous supplies to the common pool. A typical example is SmartGrids, where as well as the generators and the consumers there are local *prosumers* both consuming and contributing resources to the grid.

We note that the sequence of operation includes both *conventional* actions, like request and allocate, and *physical* actions, like contribute and appropriate. Also, it is possible for agents to 'misbehave', and not conform to the specifications, for example by appropriating more then they are allocated.

2.2. Exogenous Resource Allocation

There are many applications which require some partition of a divisible good. For example, consider a water management system with a resource (a reservoir of water) and a set of appropriators (agents) who draw water from the reservoir.

This can be formulated as a resource allocation system defined at time t by $\langle \mathcal{A}, P, m \rangle_t$, where $\mathcal{A}$ is the set of appropriators; P is the pooled resources (a divisible good); and m is the resource allocation. At each time t, m is a mapping from members of $\mathcal{A}$ to a fraction of P, $m : \mathcal{A} \mapsto [0, P]$. A *valid* allocation satisfies the constraint that $\sum_{a \in \mathcal{A}} m(a) \leqslant P$.

There are various ways of determining m_t, for example by forming a queue, by auctions [Kremer and Nyborg 2004], or cake-cutting algorithms [Brams and Taylor 1996]. Suppose it is determined by a queue, then letting $m_t(i) = r_i$ be the resources allocated to the ith agent in the queue at time t, the utility u_i of agent i is given by:

$$u_i = \begin{cases} r_i, & \text{if} \quad \sum_{j=1}^{i} r_j \leqslant P \\ 0, & \text{otherwise} \end{cases}$$

In other words, resources are allocated to the agents in the front of the queue so that it does not exceed the resources available P, and everyone else gets nothing.

However, this is the utility at one time point. Ideally each agent should maximise $\max \sum_{t=0}^{eh} u_i(t)$, where eh is some 'event horizon' at which the agent is still getting utility from the resource. A 'shortsighted' appropriator might get an optimal allocation in the short-term, but in doing so, deplete the resource and get 0 thereafter; and in fact a strategy that yields sub-optimal resources in the short-term, and does not deplete the resource (i.e. the sum of the appropriations in any one timeslice is $\leqslant P$), might provide more resources overall. Therefore, the 'optimal' distribution of resources is less important than the endurance of the distribution mechanism (cf. [Ostrom and Hess 2006, p. 68]: "systems that ... allow for change may be suboptimal in the short run but prove wiser in the long run").

This utility assumes a valid allocation and that the agent appropriates only the resources it has been allocated. These are assumptions that can be made in centralised, fully co-operative system, but we cannot necessarily make in an open system. It is necessary to distinguish between and deal with both intentional violations as well as unintentional ones, and so we require a distributed, self-organised solution.

2.3. Endogenous Resource Allocation

CPR management by an institution with endogenous resources requires that each agent both makes provision to and appropriates resources from the common pool. The agents must comply with the rules concerning both provision and appropriation, but in an open system, this includes dealing with intentional violations as well as unintentional ones.

Analysing the problem of individual resource contribution with potential rule-violations in an endogenous CPR can be considered as a linear public good (LPG) game [Gaechter 2006]. This game has proved useful for examining the free rider hypothesis, and the incentives for voluntary contributions, in both laboratory-based simulations and agent-based modelling. In a typical LPG game, n people or agents form a group or *cluster*. All cluster members individually possess a quantity of resource. Each cluster member $i \in \{1, \ldots, n\}$ decides independently to contribute resources $r_i \in [0, 1]$ to the public good. The contributions from the whole cluster are summed and the utility (payoff) u_i for each player (agent) i is given by:

$$u_i \; = \; \frac{\alpha}{n} \sum_{j=1}^{n} r_j + \beta(1 - r_i), \quad \text{where} \quad \alpha > \beta \quad \text{and} \quad \frac{\alpha}{n} < \beta$$

The first term represents the payoff from the public good (the 'public payoff'), whereby the sum of the individual contributions are distributed equally among the n cluster members. The second term represents the payoff from the resources withheld from the public good (the 'private payoff') irrespective of how much was contributed individually and collectively. The coefficients α and β represent the relative weight of the public/private payoffs respectively. If the above conditions on α and β hold, a rational but selfish agent has an incentive to contribute 0 to the public good, i.e. to free ride, so that:

- The dominant strategy is defect: the individual allocation is greatest when a member contributes 0 and every other cluster member contributes 1;
- The collective payoff is least when every cluster member contributes 0, but increases as contributions increase;
- The collective payoff is greatest when all cluster members contribute fully.

3. BACKGROUND

This section reviews the background work relevant to addressing the problem posed in the previous section. This includes the institutional approach to CPR management of Ostrom [1990], the concept of institutionalised power [Jones and Sergot 1996], and the dynamic specification of norm-governed systems [Artikis 2011].

3.1. Self-Governing the Commons

Based on extensive fieldwork, Ostrom [1990] argued that management of common-pool resources (CPRs) need not lead to a 'tragedy of the commons' as predicted by game theory, the tragedy being that a group of self-interested, autonomous and rational actors required to share a common but limited resource will inevitably act in the short-term in such a way as to deplete the resource, even if that is in none of their interests in the long-term. Ostrom showed that in spite of this result, there was an alternative to privatisation or centralised control of the resource. She observed that in many cases, for example in Spain, Switzerland, Japan and the US, communities were able to manage their own affairs by defining *institutions* to govern their commons.

Ostrom observed that common-pool resource (CPR) management problems in human societies have often been resolved through the 'evolution' of institutions. Ostrom defined an institution as a "set of working rules that are used to determine who is eligible to make decisions in some arena, what actions are allowed or constrained, ... [and] contain prescriptions that forbid, permit or require some action or outcome" [Ostrom 1990, p. 51]. She also maintained that the rule-sets were conventionally agreed (ideally by those affected by them), mutually understood, monitored and enforced; that they were nested; and that they were mutable.

On the issue of nesting, Ostrom [1990, p. 52] distinguished three levels of rules. These were, at the lowest level, *operational-choice* rules, which were concerned with the processes of resource appropriation, provision, monitoring and enforcement. At the middle level, *collective-choice* rules were concerned with selecting the operational rules, as well as processes of policy-making, role assignment and dispute resolution. At the highest level, *constitutional-choice* rules indirectly affected the operational rules by determining who is eligible to, and what specific rules are to be used to, define the set of collective-choice rules.

The nesting of rules was important for the process of *institutional change* for two reasons. Firstly, the changes which constrain action at a lower level occur in the context of an apparently 'fixed' set of rules at a higher level, i.e. what Ostrom referred to as a *decision arena*. These rules could yet be changed in another decision arena at a higher level, and so on. Secondly, lower level rules were easier and less 'costly' to change than the higher level rules. This nesting, it was found, increased the stability of strategies and expectations of those individuals having to interact with others in the context of the institutional setting.

However, Ostrom also observed that there were occasions when the institutions were *enduring*, and others where they were not. Accordingly, eight design principles were identified for *self*-management of common-pool resources (CPRs) to endure [Ostrom 1990, p. 90]. These are shown in Table I.

Table I: Ostrom's Principles for Enduring Institutions.

1 Clearly defined boundaries: those who have rights or entitlement to appropriate resources from the CPR are clearly defined, as are its boundaries.

2 Congruence between appropriation and provision rules and the state of the prevailing local environment.

3 Collective-choice arrangements: in particular, those affected by the operational rules participate in the selection and modification of those rules.

4 Monitoring, of both state conditions and appropriator behaviour, is by appointed agencies, who are either accountable to the resource appropriators or are appropriators themselves.

5 A flexible scale of graduated sanctions for resource appropriators who violate communal rules.

6 Access to fast, cheap conflict-resolution mechanisms.

7 Existence of and control over their own institutions is not challenged by external authorities.

8 Systems of systems: layered or encapsulated CPRs, with local CPRs at the base level.

3.2. Institutionalised Power

Following the third principle, if the set of working rules defining an institution contains "prescriptions that forbid, permit or require some action or outcome", and specifies formally "who is eligible to make decisions", it is generally not a specific agent that is eligible to make decisions, but instead it is an agent that occupies a designated role, that is *empowered* to make those decisions.

Therefore, we need to represent the concepts of role, role assignment [Sandhu et al. 2000] and *institutionalised power* [Jones and Sergot 1996]. The term institutionalised power refers to that characteristic feature of institutions, whereby designated agents, often acting in specific roles, are empowered to create or modify facts of special significance in that institution (*institutional facts*), through the performance of a designated action, e.g. a speech act.

This necessitates defining a role-assignment protocol that appoints a specific agent to a role. It must also be possible to change which agent occupies that role, for example if the appointed agent leaves the system, performs badly or incorrectly, or is unable to execute the duties associated with the role. To deal with assignment and change, we need dynamic specification of norm-governed systems.

3.3. Dynamic Specification of Norm-Governed Systems

Artikis [2011] defined a framework that allows agents to modify the rules or protocols of a norm-governed system at runtime. This framework defines three components: a specification of a norm-governed system, a protocol-stack for defining how to change the specification, and a topological space for expressing the 'distance' between one specification instance and another.

The study of legal, social and organisational systems has often been formalised in terms of norm-governed systems. The framework maintains the standard and long established distinction between physical capability, institutionalised power, and permission (see e.g. [Jones and Sergot 1996] for illustrations of this distinction). Accordingly, a specification of a norm-governed multi-agent system expresses five aspects of social constraint: the physical capabilities; the institutionalised powers; the permissions, prohibitions and obligations of the agents; the sanctions and enforcement policies that deal with the performance of prohibited actions and non-compliance with obligations; and the designated roles of empowered agents.

Underpinning this specification is a communication language. This language is used to define a set of protocols for conducting the business of the institution. In the framework, the protocol *stack* is used by the agents to modify the rules or protocols of a norm-governed system at runtime. This stack defines a set of object level protocols, and assumes that during the execution of an object protocol the participants could start a meta-protocol to (try to) modify the object-level protocol. The participants of the meta-protocol could initiate a meta-meta protocol to modify the rules of the meta-protocol, and so on. In addition to object- and meta-protocols, there are also 'transition' protocols. These protocols define the conditions in which an agent may initiate a meta-protocol, who occupies which role in the meta-protocol, and what elements (the *degrees of freedom*: DoF) of an object protocol can be modified as a result of the meta-protocol execution.

For example, for the first principle, we need to define who is, and who is not, a *member* of an institution. An agent can join the institution if they satisfy certain criteria, and can be excluded if they do not comply to the rules (more precisely, some other agent is empowered to admit or exclude them, under the respective conditions). To do this, we specify two types of method, one for joining and another for exclusion. Joining is a form of access control by role assignment. The type of access control method is $acMethod$, which can (again, for example) be either *attribute-based*, whereby if the applicant satisfies certain qualification criteria then it is automatically admitted, or *discretionary*, i.e. an applicant must satisfy another agent's criteria, who is acting on behalf of the institution in its appointed role. The type of exclusion method is $exMethod$, which can be either by *jury*, in which case the institution members vote on whether or not to exclude a non-complying agent, or again *discretionary*, i.e. some specific agent decides whether or not to exclude an agent.

Each type of method is a DoF, and with two values for each method, this gives four possible *specification instances*. This is the basis for defining a *specification space* as a 2-tuple, where one component is the set of all possible specification instances and the other component is a function d which defines a 'distance' between any pair of elements in the set. Note that we can imagine more access control and exclusion mechanisms, so more specification instances, and so a larger specification space.

4. OSTROM INSTITUTIONS AS DYNAMIC SPECIFICATIONS

In this section, we describe a methodology for sociologically-inspired computing, and apply it to cast Ostrom's definition of an institution (Section 3.1) as a dynamic specification of a norm-governed system (Section 3.3). From this we derive a formal model of a multi-agent system to address the problem of exogenous resource allocation.

4.1. Methodology

A methodology for sociologically-inspired computing is illustrated in Figure 1. We are here building on the synthetic method underlying research in artificial societies and artificial life [Steels and Brooks 1994], and echoing other attempts to apply ideas from the social sciences to the design of computational systems [Edmonds et al. 2005]. The method also has much in common with the CosMos methodology [Andrews et al. 2010] for biologically-inpsired computing, which identifies three phases of discovery, development and exploration; these correspond to our steps of formal characterisation, principled operationalisation, and controlled experimentation. We stress artificial experimentation because in these experiments we investigate artificial societies, rather then try to model human societies using real-world data.

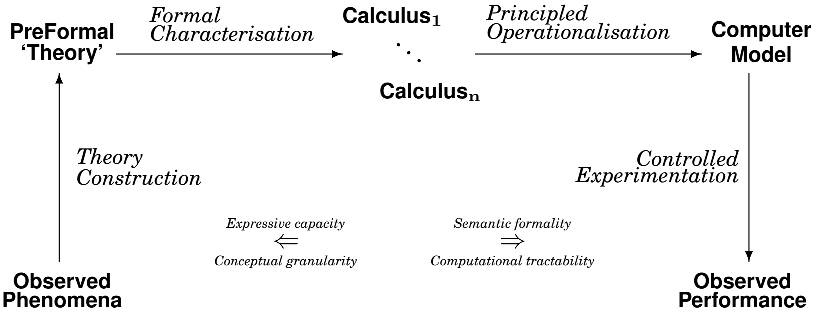

Fig. 1: Sociologically-Inspired Computing

We start from an observed phenomenon, for example a human social, legal or organisational system. The process of *theory construction* creates a pre-formal 'theory', usually specified in a natural language. Ostrom [1990] comes into this category, as it is an evidence-based theory of enduring institutions, with some notational formalism, and primarily natural language descriptions. The process of *formal characterisation* represents such theories in a calculus of some kind, where by calculus we mean any system of calculation or computation that is based on symbolic representation and manipulation. This representation can be at different levels of abstraction depending on the intended role of the calculus: expressive capacity or conceptual granularity with regard to 'theory'; computational tractability or declarative semantics with regard to implementation. The step of *principled operationalisation* embeds such formal representations in simulations which can include detailed implementation of individual agents. Finally, the computer model can be animated or executed and the performance of the model can be observed.

4.2. Institutional Rules and the Protocol Stack

From Section 3.3, the three elements of Artikis' framework [Artikis 2011] were a specification of a norm-governed system, a protocol stack, and a specification space.

Firstly, the institutional rules of Ostrom are characterised as a norm-governed system. As such, the specification will define the previously mentioned social constraints and aspects of institutional action: the physical capabilities, institutionalised powers, permissions, prohibitions and obligations of the agents; the sanctions and enforcement policies that deal with the performance of prohibited actions and non-compliance with obligations; and the designated roles of empowered agents. The Event Calculus (EC) [Kowalski and Sergot 1986] is used as the calculus for formal characterisation.

Secondly, the nesting of operational-choice rules within (social) collective-choice rules within constitutional-choice rules is treated by the object, meta- and meta-meta-protocols, and we handle *institutional change* within the framework of dynamic specifications. This proposal is illustrated in Figure 2. We show the type of rule in Ostrom's framework on the left, and the protocol that we will specify in Artikis' framework on the right. For example, the appropriation and provision operational-choice rules of Ostrom are implemented by actions in an object-level protocol for the LPG game; similarly the monitoring and enforcement rules are implemented by protocols for access control and exclusion. At the meta-level, there are protocols which change object level rules, i.e. through role assignment and choosing the DoF values for the access control and exclusion methods.

For the current work, we do not define any constitutional-choice rules, although we will observe that there are several layers of the protocol stack which contain modifiable social collective-choice rules.

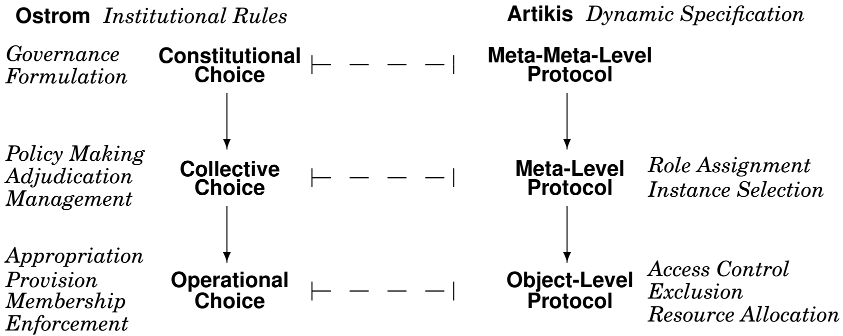

Fig. 2: Institutional Rules as a Protocol Stack

Artikis' framework originally defined the specification space as a metric space. Instead of a metric space, we represent the set of specification instances as nodes on a graph with a constant 'distance' d between any two nodes. Note that for the simple specification space here we assume a fully-connected graph, but that does not generally hold: there could be some specification instances which are not 'reachable' from others, and some which may not reachable at all. There may even be normative constraints on the transition: the institutions's members may not be empowered or permitted to make certain changes. However, we will not attach constraints or a value to d in the subsequent development; however, we find it convenient to retain it for further work in representing the 'transaction cost' of modifying operational-, collective- and constitutional-choice rules, as discussed in Section 6.

4.3. Institutional Rules and Roles

Following Ostrom [1990, pp. 52–53] as above, the institutional rules are divided into three types, ***OC***, ***SC*** and ***CC***, where OC = *operational-choice* rules, SC = *(social) collective-choice* rules, and CC = *constitutional-choice* rules.

To perform the resource allocation, we need to identify the institutional *rules* for the resource allocation itself, and the institutional *roles* whereby empowered agents perform conventional actions with institutional significance.

We identify four roles: *member*, which is the standard role for membership of an institution in order to participate in the resource allocation process; *gatekeeper*, which is empowered to assign the role of *member*; *monitor*, which is empowered to remove the role of *member*, and *head*, which is empowered to assign to the *gatekeeper* and *monitor* roles and to perform the resource allocation according to the chosen operational *method*.

We define four types of method. These are $acMethod$, $exMethod$, $raMethod$ and $wdMethod$. As above, $acMethod$ is the type of access control method: intuitively, the idea is that non-member agents will apply to join the institution, and the *gatekeeper* agent applies the access control method to determine admission or not. $exMethod$ is the type of exclusion method, correspondingly used by the *monitor* to remove membership of the institution. $raMethod$ is the type of resource allocation method, e.g. largest first, smallest first, in turn, priority, ration, etc., and used by the *head* to determine the resource allocation mapping. $wdMethod$ is the type of winner determination method used in determining collective-choices by a vote, e.g. plurality, runoff, borda, etc. [Tideman 2006]. The winner of a vote will be determined by the appropriate winner determination method.

Letting $\mathcal{M}$ be the set of members of the institution, $\{v_a(\cdot)\}_{a \in \mathcal{M}}$ denote a set of expressed preferences on an issue by each agent $a \in \mathcal{M}$, and k be a fixed winner determination method, then the set of nested institutional rules for resource allocation is illustrated in Figure 3.

$$\begin{tikzcd}
(head)\ scr_1 : \{v_a(\cdot)\}_{a \in \mathcal{M}} \times k \to wdMethod \arrow[drr] & & \\
& & (head)\ scr_2 : \{v_a(\cdot)\}_{a \in \mathcal{M}} \times wdMethod \to raMethod \arrow[dl] \\
& (head)\ ocr_1 : (\mathcal{M} \mapsto [0, P]) \times raMethod \to (\mathcal{M} \mapsto [0, P]) &
\end{tikzcd}$$

Fig. 3: Nesting of *OC* and *SC* rules for resource allocation

In Figure 3, the social collective-choice rule $scr_1 \in SC$ maps a set of expressed preferences to a winner determination method according to k; another social collective-choice rule $scr_2 \in SC$ maps a set of expressed preferences to a resource allocation method according to this winner determination method; and an operational-choice rule $ocr_1 \in OC$ maps a set of demands (i.e. a mapping $\mathcal{M} \mapsto [0, P]$) to a set of allocations (also $\mathcal{M} \mapsto [0, P]$) according to this resource allocation method. Thus ocr_1 is the resource allocation rule for the institution, i.e. it computes m_t as specified in Section 2.2.

Similarly, in Figure 4, the social collective-choice rule $scr_3 \in SC$ maps a set of expressed preferences on a social collective-choice rule $scr_i \in SC$ ($i \in \{4, 5\}$) to a winner determination method according to k. $scr_4 \in SC$ is the *gatekeeper* role assignment rule, and maps a set of expressed preferences to a designated member of $\mathcal{M}$, i.e. the *gatekeeper*, according to its winner determination method. $scr_5 \in SC$ maps a set of expressed preferences to an access control method according to its winner determination method; and an operational-choice rule $ocr_2 \in OC$ maps an application from an agent not in $\mathcal{M}$ to a boolean outcome depending on the selected access control method. From the perspective of role assignment, ocr_2 is the *member* role assignment rule, and scr_4 is the *gatekeeper* role assignment rule.

Finally, note that we have assumed the role of *head*, and certain winner determination methods k to be fixed. This is in keeping with Ostrom's idea of decision arenas, in which changes to some rules are made against a backdrop of other, fixed, rules. In practice, however, both of these are degrees of freedom, although we will not treat them as such in the dynamic specification of the institutional rules.

$$\begin{tikzcd}
 & (head)\ scr_3 : \{v_a(\cdot)\}_{a \in \mathcal{M}} \times SC \times k \rightarrow wdMethod \arrow[dl] \arrow[dr] & \\
(head)\ scr_4 : \{v_a(\cdot)\}_{a \in \mathcal{M}} \times wdMethod \rightarrow \mathcal{M} \arrow[dr] & & (head)\ scr_5 : \{v_a(\cdot)\}_{a \in \mathcal{M}} \times wdMethod \rightarrow acMethod \arrow[dl] \\
 & (gatekeeper)\ ocr_2 : \mathcal{M}^c \times acMethod \rightarrow Bool &
\end{tikzcd}$$

Fig. 4: Nesting of *OC* and *SC* rules for role assignment

4.4. Formal Model

We will now instantiate a formal model of an institution for exogenous resource allocation. The intuitive idea is that the agents will form into clusters (institutions). At each time-point t, the agents in $\mathcal{A}$ will manage membership using the access control and exclusion methods, and each cluster will determine a resource allocation method congruent with the state of the environment, using a specification instance l of the dynamic specification $\mathcal{L}$.

Let $\mathcal{IC}_t$ be a multi-agent system at time t defined by:

$$\mathcal{IC}_t = \langle \mathcal{A}, \mathcal{I}, \mathcal{L}, d \rangle_t$$

where (omitting the subscript t if clear from context):

- $\mathcal{A}$ is the set of all agents;
- $\mathcal{I}$ is the set of institutional clusters;
- $\mathcal{L}$ is a dynamic norm-governed system specification, defining a specification space;
- d is a distance function defined on specification instances of $\mathcal{L}$.

Each institutional cluster $I_t \in \mathcal{I}_t$ is defined by:

$$I_t = \langle \mathcal{M}, l, \epsilon \rangle_t$$

where (again omitting the subscript t if clear from context):

- $\mathcal{M}$ is the set of member agents, such that $\mathcal{M} \subseteq \mathcal{A}$
- l is a specification instance of $\mathcal{L}$; and
- ϵ is the cluster's local environment, a pair $\langle Bf, If \rangle$.

The local environment ϵ is a representation of Bf, the set of 'brute' facts whose values are determined by the physical state, including the value of the common-pool resource P; and If the set of 'institutional' facts, whose values are determined by the conventional state, including the roles assigned to members of $\mathcal{M}$, the methods for resource allocation, role assignment and winner determination, and the institutionalised powers, permissions and obligations. The values in If are also the DoF values which determine the specification instance l of the dynamic specification $\mathcal{L}$.

This, we contend, completes the formal characterisation step of the methodology, and establishes proposition **p1**, that Ostrom's institutions for common-pool resource management can be construed from the perspective of dynamic specifications for norm-governed systems. In the next section, we focus on the refinement of $\mathcal{L}$ to perform the resource allocation, for self-organisation and specifically its relation to the principles of enduring institutions identified by Ostrom.

5. THE FORMAL CHARACTERISATION OF SOCIO-ECONOMIC PRINCIPLES

In this section, we address proposition **p2**. In particular, we show that by our methodological process of formal characterisation, we can give a logical axiomatisation of the institutional rules as a dynamic specification $\mathcal{L}$ that performs the resource allocation in terms of the principles for enduring institutions. We proceed by giving a short summary of the Event Calculus (EC), an action language from Artificial Intelligence for reasoning about actions and events. We then list the EC *fluents* which describe the institutional facts If of the local environment. Using these fluents, we provide a logical axiomatisation, in EC, of six of the eight principles for enduring institutions specified by Ostrom (Principles 1–6, Table I).

Note that Principle 7 states that the existence of and control over the institutions is not challenged by external authorities. This means, external governmental officials should recognise the right to organise, so that the rules defined by the institution cannot be easily overruled by outsiders. In this work there is no simulation of an external authority, thus Principle 7 is provided *de facto*, simply by not modelling an external authority that can adversely affect the institution's right to self-organise (i.e. the rules of the institutions will not be overruled by outsiders). Principle 8 is discussed later in the Research Challenges.

5.1. The Event Calculus

To specify the axiomatisation of Ostrom's socio-economic principles of enduring institutions in the concepts of a norm-governed system, we use a language for representing and reasoning about action, agency, social constraints and change. There are various alternative languages; we use the Event Calculus (EC) [Kowalski and Sergot 1986] for clarity of exposition, for executable specification, and as a formal specification for implementation.

The EC is a logic formalism for representing and reasoning about actions or events and their effects. The EC is based on a many-sorted first-order predicate calculus. For the version used here, the underlying model of time is linear, so we use non-negative integer time-points (although this is not an EC restriction). We do not assume that time is discrete (the numbers need not correspond to a uniform duration) but we do impose a relative/partial ordering for events: for non-negative integers, $<$ is sufficient.

An *action description* in EC includes axioms that define: the action occurrences, with the use of happensAt predicates; the effects of actions, with the use of initiates and terminates predicates; and the values of the fluents, with the use of initially and holdsAt predicates. Table II summarises the main EC predicates. Variables, that start with an upper-case letter, are assumed to be universally quantified unless otherwise indicated. Predicates, function symbols and constants start with a lower-case letter.

Where F is a *fluent*, which is a property that is allowed to have different values at different points in time, the term $F = V$ denotes that fluent F has value V. Boolean fluents are a special case in which the possible values are $true$ and $false$. Informally, $F = V$ holds at a particular time-point if $F = V$ has been *initiated* by an action at some earlier time-point, and not *terminated* by another action in the meantime.

Events initiate and terminate a period of time during which a fluent holds a value continuously. Events occur at specific times (when they $happen$). A set of events, each with a given time, is called a *narrative*.

The utility of the EC comes from being able to reason with narratives. Therefore the final part of an EC specification is the domain-independent 'engine' which computes what fluents hold, i.e. have the value $true$ in the case of boolean fluents, or what value a fluent takes, for each multi-valued fluent. This can be used to compute a 'state' of the specification in terms of the fluents representing institutional facts. This state changes

Table II: Main Predicates of the Event Calculus.

Predicate	Meaning
Act happensAt T	Action Act occurs at time T
initially $F = V$	The value of fluent F is V at time 0
$F = V$ holdsAt T	The value of fluent F is V at time T
Act initiates $F = V$ at T	The occurrence of action Act at time T initiates a period of time for which the value of fluent F is V
Act terminates $F = V$ at T	The occurrence of action Act at time T terminates a period of time for which the value of fluent F is V

over time as event happen, and includes the roles, powers, permissions and obligations of agents, and the values assigned to each of the methods.

5.2. Fluents

The institutional facts in If required for the logical axiomatisation are represented as fluents $F = V$ in the EC. Those fluents in whose values we are interested are shown in Table III. For each fluent, we give its name, range, and which principle contain actions which initiate (or change) its value. Note that in the arguments, A is an agent, I is the institution, and M here is a 'motion' which requires vote in I. Below, we explain further the use of the fluents in the principles.

The first four fluents record the roles that agents occupy: note this formulation means that it is possible for an agent to occupy more than one role in a single institution, and be a member of more than one institution. The role of head is determined in a different decision arena.

The multi-valued fluent $\mathit{raMethod}$ specifies which resource allocation method the *head* should use in determining resource allocation according to Principle 2; while $\mathit{acMethod}$ and $\mathit{exMethod}$ specify respectively which access control and exclusion method the *gatekeeper* should use for *member* role assignment and exclusion for Principle 1. There is one fluent $\mathit{wdMethod}$ for the winner determination method for each of the social collective-choice rules that require a vote in Principle 3. The fluent $\mathit{adrMethod}$ specifies the alternative dispute resolution method to be used in Principle 6.

The boolean-valued fluent $\mathit{applied}$ is true if agent A has applied to join (become a member) of institution I, and is used in determining the (membership) boundaries in Principle 1. The fluent $\mathit{demanded}$ takes a value in $[0, P]$, giving the resources demanded by an agent A in the current time-slice. The fluent $\mathit{demand\_q}$ is a list, initially empty, of all the requests made by the members of the institution in the current time-slice. Both these fluents are used in Principle 2.

The fluent $\mathit{vote\_q}(M, I)$ is a list of votes on motion M, while $\mathit{voted}(M, I)$ is a sorted list of agents who have voted on M. In this way the ballot is private unless the vote actions are public. When the head calls for a ballot on motion M, its status is open, until the head closes the ballot. These fluents are required to implement the collective choice arrangements of Principle 3.

The fluent $\mathit{reported}(B, I)$ records a 2-tuple, comprising B's observation of the resource level P at a specific time, and the fluent $\mathit{monitoring\_freq}(I)$ is the frequency which agents appointed to the $\mathit{monitor}$ role should report their observations. These are used in Principle 4.

Table III: EC Fluents for Institutional Principles

Fluent (F)	**Range** (V)	**Principle** (F)
$role\_of(A, member, I)$	$boolean$	1
$role\_of(A, gatekeeper, I)$	$boolean$	1
$role\_of(A, monitor, I)$	$boolean$	4
$role\_of(A, head, I)$	$boolean$	n/a
$acMethod(I)$	$\{attribute, discretionary\}$	1
$exMethod(I)$	$\{jury, discretionary\}$	1
$raMethod(I)$	$\{ration, priority, torno, \ldots\}$	2
$wdMethod(scr_i, I)$	$\{plurality, runoff, borda, \ldots\}$	3
$adrMethod(I)$	$\{arb, med, neg, \ldots\}$	6
$applied(A, I)$	$boolean$	1
$demanded(A, I)$	$[0, P]$	2
$demand\_q(I)$	list of $(agent, [0, P])$	2
$vote\_q(M, I)$	list of $votes$	3
$voted(M, I)$	list of $agents$	3
$status(M, I)$	$\{open, closed\}$	3
$reported(B, I)$	$([0, P], time)$	4
$monitoring\_freq(I)$	$time$	4
$ex\_sanction\_level(I)$	$integer$	5, 6
$sanction\_level(A, I)$	$integer$	5, 6
$offences(A, I)$	$integer$	5, 6
$appealed(A, S, I)$	$boolean$	6
$\mathbf{pow}(Agent, Action)$	$boolean$	n/a
$\mathbf{per}(Agent, Action)$	$boolean$	n/a
$\mathbf{obl}(Agent, Action)$	$boolean$	n/a

Two further fluents record the number of offences committed by an agent and its current sanction level. Also, the institution has a sanction level associated with exclusion, $ex\_sanction\_level(I)$. These are used in Principles 5 and 6.

The final three fluents record the (institutionalised) powers, permissions and obligations of each agent.

Note that the degrees of freedom (DoF) of the specification are the five methods, the monitoring frequency and the sanction level for exclusion. (Looking ahead, these will all be properties of the institution in the testbed specification, see Figure 5.) The other fluents are institutional facts recorded in the environment ϵ.

5.3. The Logical Axiomatisation

This section gives an indicative axiomatisation of the principles. This serves both as a proof of concept that such an axiomatisation is feasible, but also serves to unify several (hitherto apparently) disparate strands of research in access control, voting, and alternative dispute resolution. We say the axiomatisation is indicative because there are alternative formalisations, some more complex than others. This complexity is important when cost is a factor in the application and/or adaptation of the rules, as shown in Section 6.

5.3.1. Principle 1: Clearly Defined Boundaries. Principle 1 states that those who have rights or entitlement to appropriate resources from the CPR are clearly defined, as are its boundaries.

There are three aspects to axiomatising this principle and defining an institution's boundaries. The first issue is to separate those who have rights and entitlements from those who do not; and the second is to express precisely what those rights and entitlements are.

The first issue can be dealt with using role-based access control (e.g. [Firozabadi and Sergot 2004]) and defining a role-assignment protocol (cf. [Artikis and Sergot 2010]), in order to distinguish between those agents in $\mathcal{A}$ that are members of the institution (i.e. the set $\mathcal{M}$) and those who are not.

An agent is empowered to apply for membership to an institution I if it is not already a member of that institution.

$$\begin{array}{l} \mathit{apply}(A, I) \quad \mathsf{initiates} \quad \mathit{applied}(A, I) = \mathit{true} \quad \mathsf{at} \quad T \quad \leftarrow \\ \qquad \mathit{role\_of}(A, \mathit{member}, I) = \mathit{false} \quad \mathsf{holdsAt} \quad T \end{array}$$

The gatekeeper agent G is empowered to admit an agent A as a *member* to the institution I, by an assign action, depending on the access control method.

$$\begin{array}{l} \mathit{assign}(G, A, \mathit{member}, I) \quad \mathsf{initiates} \quad \mathit{role\_of}(A, \mathit{member}, I) = \mathit{true} \quad \mathsf{at} \quad T \quad \leftarrow \\ \qquad \mathbf{pow}(G, \mathit{assign}(G, A, \mathit{member}, I)) = \mathit{true} \quad \mathsf{holdsAt} \quad T \\ \mathbf{pow}(G, \mathit{assign}(G, A, \mathit{member}, I)) = \mathit{true} \quad \mathsf{holdsAt} \quad T \quad \leftarrow \\ \qquad \mathit{applied}(A, I) = \mathit{true} \quad \mathsf{holdsAt} \quad T \quad \wedge \\ \qquad \mathit{acMethod}(I) = \mathit{attribute} \quad \mathsf{holdsAt} \quad T \quad \wedge \\ \qquad \mathit{role\_of}(G, \mathit{gatekeeper}, I) = \mathit{true} \quad \mathsf{holds} \quad T \quad \wedge \\ \qquad \mathit{role\_conditions}(\mathit{member}, A, I) = \mathit{true} \quad \mathsf{holdsAt} \quad T \\ \mathbf{pow}(G, \mathit{assign}(G, A, \mathit{member}, I)) = \mathit{true} \quad \mathsf{holdsAt} \quad T \quad \leftarrow \\ \qquad \mathit{applied}(A, I) = \mathit{true} \quad \mathsf{holdsAt} \quad T \quad \wedge \\ \qquad \mathit{acMethod}(I) = \mathit{discretionary} \quad \mathsf{holdsAt} \quad T \quad \wedge \\ \qquad \mathit{role\_of}(G, \mathit{gatekeeper}, I) = \mathit{true} \quad \mathsf{holdsAt} \quad T \end{array}$$

If the $\mathit{acMethod}$ is *attribute*, then the gatekeeper is empowered to assign the role *member* provided the applicant satisfies certain (external) role conditions. The conditions could include, for example, not exceeding a fixed number of non-compliant actions, a duration since the last non-compliant action, and so on.

If the $\mathit{acMethod}$ is *discretionary*, then the *gatekeeper* is empowered to assign the role without conditions, according to its (internal) decision-making, which could yet make reference to external conditions.

Similarly, the *gatekeeper* G is empowered to exclude a *member*, and is permitted to do so when a member does not comply with the rules of the institution I; for example, it appropriates more resources than it has been allocated (and see Principles 4 and 5).

If the $\mathit{exMethod}$ is discretionary, then the *monitor* agent G is empowered to exclude an agent A (or not) as it decides. If the $\mathit{exMethod}$ is jury, then the *monitor* must have called for a vote on the issue of the exclusion of A:

$$
\begin{array}{l}
\mathit{exclude}(G, A, \mathit{member}, I) \quad \textsf{initiates} \quad \mathit{role\_of}(A, \mathit{member}, I) = \mathit{false} \quad \textsf{at} \quad T \quad \leftarrow \\
\qquad \mathit{role\_of}(A, \mathit{member}, I) = \mathit{true} \quad \textsf{holdsAt} \quad T \quad \wedge \\
\qquad \mathbf{pow}(G, \mathit{exclude}(G, A, \mathit{member}, I)) = \mathit{true} \quad \textsf{holdsAt} \quad T \\
\mathbf{pow}(G, \mathit{exclude}(G, A, \mathit{member}, I)) = \mathit{true} \quad \textsf{holdsAt} \quad T \quad \leftarrow \\
\qquad \mathit{role\_of}(G, \mathit{gatekeeper}, I) = \mathit{true} \quad \textsf{holdsAt} \quad T \quad \wedge \\
\qquad \mathit{exMethod}(I) = \mathit{discretionary} \quad \textsf{holdsAt} \quad T \\
\mathbf{pow}(G, \mathit{exclude}(G, A, \mathit{member}, I)) = \mathit{true} \quad \textsf{holdsAt} \quad T \quad \leftarrow \\
\qquad \mathit{role\_of}(G, \mathit{gatekeeper}, I) = \mathit{true} \quad \textsf{holdsAt} \quad T \quad \wedge \\
\qquad \mathit{exMethod}(I) = (\mathit{jury}, \mathit{WDM}) \quad \textsf{holdsAt} \quad T \quad \wedge \\
\qquad \mathit{ballot}(\mathit{exclude}(A), I) = V \quad \textsf{holdsAt} \quad T \quad \wedge \\
\qquad \mathit{winner\_determination}(\mathit{WDM}, V, \mathit{true})
\end{array}
$$

Note that when the $\mathit{exMethod}$ is *discretionary*, the *monitor* is empowered to exclude any *member*, but if the exclusion method is *jury*, only when the result of a ballot (a vote) is in favour of exclusion. However, the *monitor* is only *permitted* to exercise its power when that member has been sanctioned, and the sanction level of the agent has reached a threshold for exclusion which is specific to the institution and its graduated sanctions.

$$
\begin{array}{l}
\mathbf{per}(G, \mathit{exclude}(G, A, \mathit{member}, I)) = \mathit{true} \quad \textsf{holdsAt} \quad T \quad \leftarrow \\
\qquad \mathit{role\_of}(G, I) = \mathit{gatekeeper} \quad \textsf{holdsAt} \quad T \quad \wedge \\
\qquad \mathit{sanction\_level}(A, I) = S \quad \textsf{holdsAt} \quad T \quad \wedge \\
\qquad \mathit{ex\_sanction\_level}(I) = S \quad \textsf{holdsAt} \quad T
\end{array}
$$

Note that similar role-assignment axioms and conditions are required for the *head* agent to appoint or remove agents, who are already members, to the roles of gatekeeper and monitor (see e.g. Principle 4). This takes place against a fixed backdrop of rules where the role of *head* has been decided. We omit these rules here, but in practice there will be a (meta) decision arena whose remit is to appoint the *head*.

5.3.2. Principle 2: Congruence of Rules and Environment. Principle 2 states that there should be a congruence between appropriation and provision rules and the state of the prevailing local environment. For this, we need to define axioms that enable changing the specification to use a different allocation method (Principle 2), and then define axioms for participatory adaptation in the collective choice arrangements (Principle 3) through voting.

In [Pitt et al. 2011a], we investigated the interleaving of rules of social order (i.e. a norm-governed system), rules of social exchange (e.g. opinion formation) [Hegselmann and Krause 2002], and rules of social computational choice [Chevaleyre et al. 2007] to balance the choice of security policy against the available energy in an ad hoc network. The issue of concern was that a ‘stronger’ security policy is more computationally intensive, and so more energy-demanding, than a weaker one. Therefore too high a security level to protect against a non-existent battery-exhaustion attack was effectively doing the job for the attacker. The requirement was to make the security level congruent to the environment, taking into account both the actual energy level and the perceived threat.

However, it wads then shown how brute facts, such as the energy level, and institutional facts, such as the security level, could be correlated by using processes of opinion formation and social choice. In a resource allocation scenario, the brute fact is the value of the common-pool resource P, and the institutional fact is the resource allocation method $raMethod$, and the aim is to find a balance between the available resources P and the method for allocating them which satisfies (in some sense) the membership and sustains the resource.

We therefore need axioms for to determine who is empowered to make demands (cf. Principle 1), axioms for the power of the *head* agent to grant allocations which are dependent on the state of the local environment, and axioms concerning the rights and entitlements of the agents. (In addition we need axioms for ensuring participation in selecting the resource allocation method so that it is congruent with the environment, but this is considered as part of the collective choice arrangements in Principle 3.)

Agents make a demand R for resources, where R is some fraction of the pooled resources P. To make the institutional fact $demanded$ true in I, an agent must be empowered, and it is empowered if it is a member of I, it has not made demand in this time-slice, and it has not been sanctioned (see Principle 5). This enforces the 'boundary' conditions from Principle 1, as no non-member or excluded member is empowered to make a demands true as an institutional fact, i.e. their demand actions are 'noise'. A demand made by an empowered agent also adds that demand to the demand queue $demand\_q$ fluent.

$$\begin{array}{l}
demand(A,R,I) \quad \mathsf{initiates} \quad demanded(A,I) = R \quad \mathsf{at} \quad T \quad \leftarrow \\
\qquad \mathbf{pow}(A, demand(A,R,I)) = true \quad \mathsf{holdsAt} \quad T \\
demand(A,R,I) \quad \mathsf{initiates} \quad demand\_q(I) = [(A,R) \mid Q] \quad \mathsf{at} \quad T \quad \leftarrow \\
\qquad demand\_q(I) = Q \quad \mathsf{holdsAt} \quad T \quad \wedge \\
\qquad \mathbf{pow}(A, demand(A,R,I)) = true \quad \mathsf{holdsAt} \quad T \\
\mathbf{pow}(A, demand(A,R,I)) = true \quad \mathsf{holdsAt} \quad T \quad \leftarrow \\
\qquad role\_of(A, member, I) = true \quad \mathsf{holdsAt} \quad T \quad \wedge \\
\qquad demanded(A,I) = 0 \quad \mathsf{holdsAt} \quad T \quad \wedge \\
\qquad sanction\_level(A,I) = 0 \quad \mathsf{holdsAt} \quad T
\end{array}$$

Now recall that allocation of resources in our example depends on resource availability. Previously we mentioned five methods for doing the allocation: largest-first, smallest-first, queue, ration, and priority. These determine the conditions on the power of the *head* to allocate resources.

$$\begin{array}{l}
\mathbf{pow}(C, allocate(C,A,R,I)) = true \quad \mathsf{holdsAt} \quad T \quad \leftarrow \\
\qquad demanded(A,I) = R \quad \mathsf{holdsAt} \quad T \quad \wedge \\
\qquad demand\_q(I) = [(A,R) \mid Q] \quad \mathsf{holdsAt} \quad T \quad \wedge \\
\qquad role\_of(C,I) = head \quad \mathsf{holdsAt} \quad T \quad \wedge \\
\qquad raMethod(I) = queue \quad \mathsf{holdsAt} \quad T \\
\mathbf{pow}(C, allocate(C,A,R',I)) = true \quad \mathsf{holdsAt} \quad T \quad \leftarrow \\
\qquad demanded(A,I) = R \quad \mathsf{holdsAt} \quad T \quad \wedge \\
\qquad demand\_q(I) = [(A,R) \mid Q] \quad \mathsf{holdsAt} \quad T \quad \wedge \\
\qquad role\_of(C,I) = head \quad \mathsf{holdsAt} \quad T \quad \wedge \\
\qquad raMethod(I) = ration(R'') \quad \mathsf{holdsAt} \quad T \quad \wedge \\
\qquad ((R > R'' \wedge R' = R'') \quad \vee \quad (R \leqslant R'' \wedge R' = R))
\end{array}$$

The last line says that either the agent demanded more than the ration R'', in which case all it should be allocated is the ration; or it demanded less than (or equal to) the ration, in which case it should be allocated what it demanded. The axioms for the other resource allocation methods are similar and are omitted.

Allocation is closely associated with the issue of rights and entitlements. It has been argued [Firozabadi and Sergot 2004] that in access control and resource allocation situations of the type being analysed here, where there may be both 'valid' and 'invalid' demands, the notions of permission and prohibition are insufficient, and a notion of *entitlement* is required.

Therefore, for an agent A that 'should be allocated' resources, there is an entitlement of A to be allocated resources. As such, there is a corresponding obligation on another agent – the one occupying the role of *head*, i.e. C – to grant empowered demands, as determined by the allocation method, i.e.:

$$\begin{aligned}
&\mathbf{obl}(C, allocate(C, A, R, I)) = true \quad \mathsf{holdsAt} \quad T \quad \leftarrow \\
&\qquad demanded(A, I) = R \quad \mathsf{holdsAt} \quad T \quad \wedge \\
&\qquad demand\_q(I) = [(A, R) \mid Q] \quad \mathsf{holdsAt} \quad T \quad \wedge \\
&\qquad role\_of(C, I) = head \quad \mathsf{holdsAt} \quad T \quad \wedge \\
&\qquad raMethod(I) = queue \quad \mathsf{holdsAt} \quad T
\end{aligned}$$

and similarly for the other appropriation rules.

5.3.3. Principle 3: Collective-Choice Arrangements. Principle 3 concerns collective-choice arrangements: in particular, those affected by the operational rules participate in the selection and modification of those rules.

For participation in the selection of rules, we need to ensure that the members of institution I are effectively enfranchised. The concept of enfranchisement can be built up from more fundamental norms for right and entitlement. In [Pitt et al. 2006], the concept of enfranchisement was formalised in two dimensions: firstly having the right to vote, and secondly having an entitlement associated with that right. Having the right to vote included having the power (being empowered) to vote. Having the associated entitlement included providing a mechanism to count the vote in accordance with the way it was cast; and an obligation on someone, occupying a designated role, to declare a 'correct' outcome (i.e. the result should be declared according to the way the votes were cast with respect to the standing rules of the election).

The right to vote, as an institutional power, is given by the axiom:

$$\begin{aligned}
&\mathbf{pow}(A, vote(A, X, M, I)) = true \quad \mathsf{holdsAt} \quad T \quad \leftarrow \\
&\qquad status(M, I) = open \quad \mathsf{holdsAt} \quad T \quad \wedge \\
&\qquad role\_of(A, member, I) = true \quad \mathsf{holdsAt} \quad T \quad \wedge \\
&\qquad voted(M, I) = L \quad \mathsf{holdsAt} \quad T \quad \wedge \\
&\qquad \mathbf{not} \quad in\_voted(A, L)
\end{aligned}$$

This states that agent A has the power to vote on issue M in institution I if three conditions are satisfied. Firstly, that the status of the issue is $open$, i.e. an appropriately empowered agent in I has called for a vote (opened a ballot) on M, which set the fluent $status(M, I)$ to $open$; and no appropriately empowered agent in I has closed the ballot, i.e. has set $status(M, I)$ to $closed$. Secondly, the agent must have the role of $member$ in I. Note that what X denotes, either yes/no, a number, a candidate list, etc., depends on the content of M (see below). Thirdly, the agent cannot already have voted on the issue (the predicate in_voted simply checks if agent A is on the list L of agents who have voted on the issue).

Designated actions, i.e. votes by member agents, can be specified to establish the necessary institutional facts which provide the mechanism to count the vote in accordance with the way it was cast:

$$\begin{array}{l}
\mathit{vote}(A,X,M,I) \quad \mathsf{initiates} \quad \mathit{vote\_q}(M,I) = [X \mid Q] \quad \mathsf{at} \quad T \quad \leftarrow \\
\qquad \mathit{vote\_q}(M,I) = Q \quad \mathsf{holdsAt} \quad T \quad \wedge \\
\qquad \mathbf{pow}(A, \mathit{vote}(A,X,M,I)) = \mathit{true} \quad \mathsf{holdsAt} \quad T \\
\mathit{vote}(A,X,M,I) \quad \mathsf{initiates} \quad \mathit{voted}(M,I) = L \quad \mathsf{at} \quad T \quad \leftarrow \\
\qquad \mathit{voted}(M,I) = L' \quad \mathsf{holdsAt} \quad T \quad \wedge \\
\qquad \mathit{sort}([A \mid L'], L) \quad \wedge \\
\qquad \mathbf{pow}(A, \mathit{vote}(A,X,M,I)) = \mathit{true} \quad \mathsf{holdsAt} \quad T
\end{array}$$

This adds the vote X to a fluent whose value is a list of votes cast $\mathit{vote\_q}$, and the voter A to a fluent whose value is a *sorted* list of voters, so the ballot is effectively private.

Finally we also have the following axiom, which is the obligation on the designated agent C occupying the role of head to declare the result in accordance with the way the votes were cast and the effective winner determination method $\mathit{wdMethod}(M,I)$ for this issue:

$$\begin{array}{l}
\mathbf{obl}(C, \mathit{declare}(C,W,M,I)) = \mathit{true} \quad \mathsf{holdsAt} \quad T \quad \leftarrow \\
\qquad \mathit{role\_of}(C, \mathit{head}, I) = \mathit{true} \quad \mathsf{holdsAt} \quad T \quad \wedge \\
\qquad \mathit{status}(M,I) = \mathit{closed} \quad \mathsf{holdsAt} \quad T \quad \wedge \\
\qquad \mathit{vote\_q}(M,I) = Q \quad \mathsf{holdsAt} \quad T \quad \wedge \\
\qquad \mathit{wdMethod}(M,I) = \mathit{WDM} \quad \mathsf{holdsAt} \quad T \quad \wedge \\
\qquad \mathit{winner\_determination}(\mathit{WDM}, Q, W)
\end{array}$$

This axiomatises the aspect of entitlement which states that the outcome of the vote should be declared correctly, by putting an obligation on the agent occupying the role of head to declare the result depending on the votes cast on issue M given by the value of the fluent $\mathit{vote\_q}(M,I)$, and the winner determination method for issue M given by the value of the fluent $\mathit{wdMethod}(M,I)$.

For the selection of a rule affecting the participants in I, we can for example arrange for a vote on the appropriation rule by calling for a vote. Suppose that in institution i the agent in the role of *head* is c, the current allocation method is *queue*, the winner determination method for choosing the resource allocation method is *plurality*, and there are agents a, b and c. Then suppose we had the following narrative:

$$\begin{array}{l}
\mathit{open\_ballot}(c, \mathit{raMethod}, i) \quad \mathsf{happensAt} \quad 1 \\
\mathit{vote}(a, \mathit{ration}, \mathit{raMethod}, i) \quad \mathsf{happensAt} \quad 2 \\
\mathit{vote}(b, \mathit{ration}, \mathit{raMethod}, i) \quad \mathsf{happensAt} \quad 3 \\
\mathit{close\_ballot}(c, \mathit{raMethod}, i) \quad \mathsf{happensAt} \quad 4
\end{array}$$

In this narrative of four event, the head c opens a ballot on the issue of the resource allocation method $\mathit{raMethod}$. Agents a and b cast their votes, which assuming they are both empowered as members of the institution, are votes for ration to be the resource allocation method. Finally the head c closes the ballot. Then there should be some event in the narrative (at time-point 5, say):

$$\mathit{declare}(c, \mathit{ration}, \mathit{raMethod}, i) \quad \mathsf{happensAt} \quad 5$$

which changes the fluent for the corresponding issue M (in this case $raMethod$) via the axiom:

$$\begin{array}{l} declare(C, W, M, I) \quad \mathsf{initiates} \quad M(I) = W \quad \mathsf{holdsAt} \quad T \quad \leftarrow \\ \qquad \mathbf{pow}(C, declare(C, W, M, I)) = true \quad \mathsf{holdsAt} \quad T \end{array}$$

so that the appropriation rule in I is now $ration$.

For the modification of a rule affecting the participants in the institution, note that we can specify exactly the same process, but with the issue M being the winner determination rule for the appropriation rule. For more on this style of hierarchical dynamic specification, see [Artikis 2011]. For a full formalisation of a voting protocol, based on the proscriptions in Robert's Rules of Order [Robert et al. 2000], which addresses specific issues enforcing one-member-one-vote, private ballots, chair's casting vote, etc., see [Pitt et al. 2006].

5.3.4. Principle 4: Monitoring.

Principle 4 is concerned with ensuring that monitoring, of both environmental conditions and appropriator behaviour, is by appointed agencies, who are accountable to the resource appropriators or are appropriators themselves.

Logic in general and the Event Calculus in particular has a well-established usage in event recognition and environment monitoring (cf. [Dousson 1996; Domingos and Lowd 2009]). To encapsulate the principle, the EC axioms need to ensure that the monitors are appropriators themselves. To do this, we define a new role, *monitor*. We ensure that the role is occupied by an appropriator by requiring that the corresponding role assignment protocol contains the condition that the appointee be a *member* of I. The *head* is empowered to perform the assignment.

$$\begin{array}{l} assign(C, B, I) \quad \mathsf{initiates} \quad role\_of(B, monitor, I) = true \quad \mathsf{at} \quad T \quad \leftarrow \\ \qquad \mathbf{pow}(C, assign(C, B, I)) = true \quad \mathsf{holdsAt} \quad T \\ \mathbf{pow}(C, assign(C, B, I)) = true \quad \mathsf{holdsAt} \quad T \quad \leftarrow \\ \qquad role\_of(B, member, I) = true \quad \mathsf{holdsAt} \quad T \quad \wedge \\ \qquad role\_of(C, head, I) = true \quad \mathsf{holdsAt} \quad T \end{array}$$

Appointment to the *monitor* role is associated with two responsibilities. The first is the *obligation* to sample the state of the environment (i.e. the resource level P), and participate in the process of opinion exchange, as it requires monitoring of the environment and the reporting of brute facts. This is used to trigger a change to the appropriation rule congruent to the state of the environment (Principle 2) using the collective-choice protocols (Principle 3).

$$\begin{array}{l} report(B, P, I) \quad \mathsf{initiates} \quad reported(B, I) = (P, T) \quad \mathsf{at} \quad T \quad \leftarrow \\ \qquad \mathbf{pow}(B, report(B, P, I)) = true \quad \mathsf{holdsAt} \quad T \\ \mathbf{pow}(B, report(B, P, I)) = true \quad \mathsf{holdsAt} \quad T \quad \leftarrow \\ \qquad role\_of(B, monitor, I) = true \quad \mathsf{holdsAt} \quad T \\ \mathbf{obl}(B, report(B, \_, I)) = true \quad \mathsf{holdsAt} \quad T \quad \leftarrow \\ \qquad role\_of(B, monitor, I) = true \quad \mathsf{holdsAt} \quad T \\ \qquad reported(B, I) = (\_, T') \quad \mathsf{holdsAt} \quad T \\ \qquad monitoring\_frequency(I) = F \quad \mathsf{holdsAt} \quad T \quad \wedge \\ \qquad T' < T + F \end{array}$$

The second responsibility is to observe appropriations, and report this information to the *head*. The observation of appropriations is required to ensure the rules are being followed, but is not part of the narrative of events. However, it is possible for the report of a misappropriation, following an observation, to lead to a sanction (Principle 5) and even to a dispute (Principle 6). Note that the role of *monitor* empowers one agent to report another:

$$\begin{aligned}&\mathbf{pow}(B, \mathit{report}(B,A,R,I)) = \mathit{true} \quad \mathsf{holdsAt} \quad T \quad \leftarrow\\&\qquad \mathit{role\_of}(B, \mathit{monitor}, I) = \mathit{true} \quad \mathsf{holdsAt} \quad T \quad \wedge\\&\qquad \mathit{role\_of}(A, \mathit{member}, I) = \mathit{true} \quad \mathsf{holdsAt} \quad T\end{aligned}$$

The fluent that a report event initiates is discussed in the next section.

5.3.5. Principle 5: Graduated Sanctions. Principle 5 states that there should be a flexible scale of graduated sanctions for resource appropriators who violate communal rules. For example, for a first offence the sanction level is increased to 1 and the power to demand is temporarily withdrawn; for a second offence the sanction level is increased to 2 and the agent may be excluded from the institution, if the value of the fluent $\mathit{ex\_sanction\_level}$ is 2.

For example, consider the following narrative, with *member* agent a, *monitor* agent b, *head* agent c, and the *ration* appropriation rule in force, and that $r > r'$:

$$\begin{aligned}&\mathit{demand}(a,r,i) \quad \mathsf{happensAt} \quad 14\\&\mathit{allocate}(c,a,r',i) \quad \mathsf{happensAt} \quad 15\\&\mathit{appropriate}(a,r,i) \quad \mathsf{happensAt} \quad 16\\&\mathit{report}(b,a,r,i) \quad \mathsf{happensAt} \quad 17\end{aligned}$$

Agent a has violated the communal rule by appropriating resources to which it was not entitled, and is reported by monitor agent b (see Principle 4).

We can add an axiom that counts rule violation offences:

$$\begin{aligned}&\mathit{report}(B,A,R,I) \quad \mathsf{initiates} \quad \mathit{offences}(A,I) = O1 \quad \mathsf{at} \quad T \quad \leftarrow\\&\qquad \mathbf{pow}(B, \mathit{report}(B,A,R,I)) = \mathit{true} \quad \mathsf{holdsAt} \quad T \quad \wedge\\&\qquad \mathit{offences}(A,I) = O \quad \mathsf{holdsAt} \quad T \quad \wedge\\&\qquad O1 = O + 1 \quad \wedge\\&\qquad \mathit{raMethod}(I) = \mathit{ration}(R') \quad \mathsf{holdsAt} \quad T \quad \wedge\\&\qquad \mathit{allocated}(A,R',I) = \mathit{true} \quad \mathsf{holdsAt} \quad T \quad \wedge\\&\qquad R > R'\end{aligned}$$

and empower the head agent to sanction offences:

$$\begin{aligned}&\mathit{sanction}(C,A,S,I) \quad \mathsf{initiates} \quad \mathit{sanction\_level}(A,I) = S \quad \mathsf{at} \quad T \quad \leftarrow\\&\qquad \mathbf{pow}(C, \mathit{sanction}(C,A,S,I)) = \mathit{true} \quad \mathsf{holdsAt} \quad T\\&\mathbf{pow}(C, \mathit{sanction}(C,A,S,I)) = \mathit{true} \quad \mathsf{holdsAt} \quad T \quad \leftarrow\\&\qquad \mathit{role\_of}(C, \mathit{head}, I) = \mathit{true} \quad \mathsf{holdsAt} \quad T \quad \wedge\\&\qquad \mathit{offences}(A,I) = S \quad \mathsf{holdsAt} \quad T\end{aligned}$$

If an agent A is sanctioned at level 1 for a first offence, then it is not empowered to demand (as specified by Principle 2). The head agent is empowered to 'reset' the sanction level $\mathit{sanction\_level}(A,I) = 0$, so that A once again has its power, but the number of offences does not decrease (i.e. $\mathit{offences}(A,I) = 1$ still holds). If agent A violates the appropriation rule again, and is sanctioned a second time, the head agent is permitted to exclude agent A because $\mathit{sanction\_level}(A,I) = 2$ (assuming $\mathit{ex\_sanction\_level} = 2 \quad \mathsf{holdsAt} \quad T$, as specified by Principle 1). Note in fact that the number of offences, the number of graduations, and the exclusion sanction level, are all technically DoF of the specification, and we can accordingly have more or less tolerant institutions.

Graduated sanctions interleave closely with the conflict-resolution mechanisms of Principle 6, which can help treat intentional and unintentional violations differently.

5.3.6. Principle 6: Conflict Resolution. Principle 6 states that the institution should provide rapid access to low-cost conflict-resolution mechanisms, such as Alternative Dispute Resolution (ADR). ADR has numerous benefits as an alternative to litigation, including lower cost, shorter time, and damage limitation. It can preserve and even strengthen relationships among the parties.

The rise in importance of ADR methods is due to the numerous benefits offered to the parties involved in a dispute, coupled with the well known shortcomings of litigation. For example, very few lawsuits filed actually go to trial, and of this an even smaller proportion arrive at a verdict. This is often due to a settlement being reached just prior to the end, or the case breaking down leading to a retrial. In both cases, the process may have taken some time, and both parties will have gone to considerable unnecessary expense.

The benefits of ADR include lower cost and relative speed. It also allows the parties involved to have more control over their dispute, and so settlement, as they chose the procedure and terms and conditions, and any third party required can be determined by those involved. Processes such as mediation and arbitration provide the parties with an opportunity for greater control over the dispute resolution process, and allow them to resolve their conflict in a more creative way than might be possible if it were left to a decision by a judge or jury.

The axiomatisation of ADR is therefore a key element of providing low-cost, rapid conflict-resolution mechanisms for self-governing commons (cf. [Katsh 2006]. In this work we will only present a simple appeals procedure, although a more refined approach can be defined.

From the specification of the previous two principles, once the monitor has reported an agent, its number of offences is incremented. Given a certain number of offences, the head is empowered to apply a sanction. However, the sanctioned agent is empowered to make an appeal against a sanction:

$$\begin{array}{l}
\mathit{appeal}(A,S,I) \quad \textsf{initiates} \quad \mathit{appealed}(A,S,I) = \mathit{true} \quad \textsf{at} \quad T \quad \leftarrow \\
\qquad \mathbf{pow}(A, \mathit{appeal}(A,S,I)) = \mathit{true} \quad \textsf{holdsAt} \quad T \\
\mathbf{pow}(A, \mathit{appeal}(A,S,I)) = \mathit{true} \quad \textsf{holdsAt} \quad T \quad \leftarrow \\
\qquad \mathit{role\_of}(A, \mathit{member}, I) = \mathit{true} \quad \textsf{holdsAt} \quad T \quad \wedge \\
\qquad \mathit{sanction\_level}(A,I) = S \quad \textsf{holdsAt} \quad T
\end{array}$$

An agent C is empowered to uphold or reject the appeal, if it occupies the role of head, the ADR method is arbitration (arb; note we assume the head agent is by default the arbiter), and the (allegedly offending) agent A has in fact appealed. Upholding the appeal removes the sanction and decrements the offence count, rejecting it (specification omitted here) does neither.

$$\begin{array}{l}
\mathit{uphold}(C,A,S,I) \quad \textsf{initiates} \quad \mathit{sanction\_level}(A,I) = \mathit{S1} \quad \textsf{at} \quad T \quad \leftarrow \\
\qquad \mathbf{pow}(C, \mathit{uphold}(C,A,S,I)) = \mathit{true} \quad \textsf{holdsAt} \quad T \ \wedge \\
\qquad \mathit{sanction\_level}(A,I) = S \quad \textsf{holdsAt} \quad T \quad \wedge \\
\qquad S1 = S - 1
\end{array}$$

$$\begin{array}{l}
\mathit{uphold}(C,A,S,I) \quad \textsf{initiates} \quad \mathit{offences}(A,I)=O1 \quad \textsf{at} \quad T \quad \leftarrow \\
\qquad \mathbf{pow}(C,\mathit{uphold}(C,A,S,I))=\mathit{true} \quad \textsf{holdsAt} \quad T \quad \wedge \\
\qquad \mathit{offences}(A,I)=O \quad \textsf{holdsAt} \quad T \quad \wedge \\
\qquad O1=O-1 \\
\mathbf{pow}(C,\mathit{uphold}(C,A,S,I))=\mathit{true} \quad \textsf{holdsAt} \quad T \quad \leftarrow \\
\qquad \mathit{role\_of}(C,\mathit{head},I)=\mathit{true} \quad \textsf{holdsAt} \quad T \quad \wedge \\
\qquad \mathit{adrMethod}(I)=\mathit{arb} \quad \textsf{holdsAt} \quad T \quad \wedge \\
\qquad \mathit{appealed}(A,S,I)=\mathit{true} \quad \textsf{holdsAt} \quad T
\end{array}$$

This simple appeals procedure allows the head agent, acting as the arbiter, to apply some form of 'common sense' reasoning to the application of the graduated sanctions. Note that we are here assuming that the monitor is a completely reliable observer and reporter. A more substantive appeals procedure would take into account that the monitor may incorrectly report an agent's appropriation, so it would need to give grounds for an appeal, and that might require a more complex protocol [Katsh 2006].

6. EXPERIMENTAL EVALUATION

This section describes the testbed that has been developed to experiment with six of Ostrom's design principles for common-pool resources [Ostrom 1990, p. 90], using exogenous resource provision. We aim to show that the axiomatisation of Section 5.3 allows us to test whether these principles are necessary and sufficient conditions for enduring electronic institutions (**p3**).

In this testbed, the EC narrative is generated and assimilated strictly in the order in which events happen, and there are no persistence disturbing effects between events affecting the same fluent. We also treat all the fluents (even boolean-valued fluents) as multi-valued fluents whose values are changed with initiates rather than terminates. Under these circumstances, as in [Farrell et al. 2005], the axioms of the specification can be implemented in C++ as state-transition constraints and, if satisfied, update the initiated fluent as part of a global state against which the next event can be evaluated.

The next section describes the testbed's classes, states and algorithm, followed by a specification of the parameters that were used in the experiments. The experimental results obtained by selecting subsets of Ostrom's design principles are presented and interpreted, and we conclude with an evaluation of these results and the limitations of the testbed.

6.1. Testbed Specification

6.1.1. Classes. The diagram in Figure 5 shows the relationship between the classes of the testbed. Note that the actions and fluents of the previous section are all methods or properties of a class.

The classes include the different roles that an agent can (adopt or) be assigned, like *member* or *non-member*, *head*, *monitor* and *gatekeeper*. The roles member (of an institution) and non-member are mutually exclusive and determined at the start of each simulation via $\mathit{acMethod}$. This allows for a clear boundary of agents who are permitted to appropriate from the CPR [Cox et al. 2010]. It is a restriction of the institution that agents should be members in order to be assigned the role of head, gatekeeper or monitor. The head is empowered to assign a single agent to more than one role.

Every agent has a name ($\mathit{ag\_name}$), an $\mathit{activity}$ status, saying whether it is permitted to appropriate resources or not, and a degree of compliance ($\mathit{compliancy\_degree}$), meaning to what extent an agent complies to the rules, especially the amount of resource it is allowed to appropriate. Any agent (member or non-member) has the physical ca-

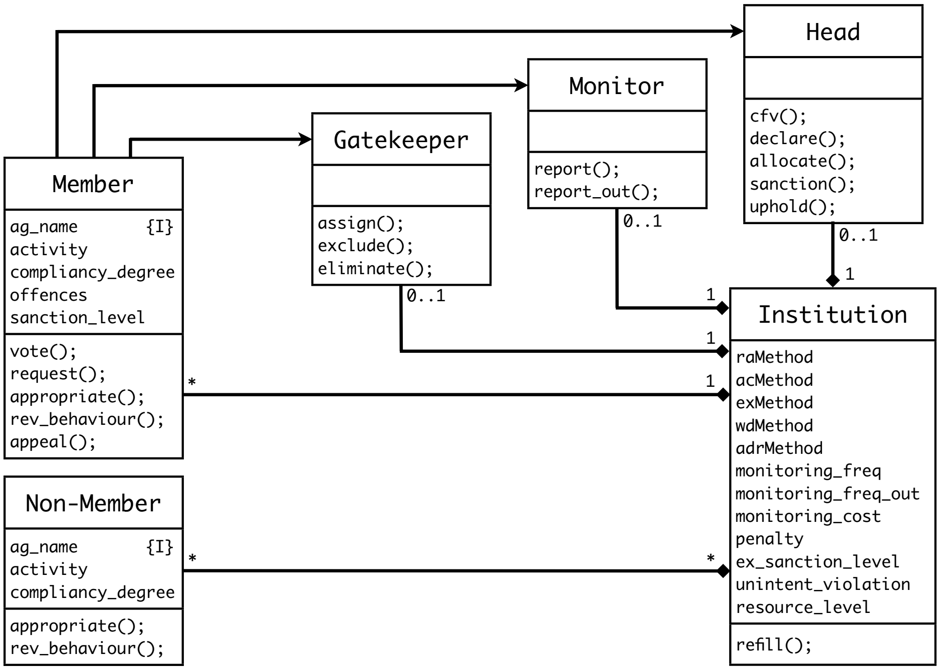

Fig. 5: Testbed class diagram

pability to *appropriate* resources, but only members are permitted to exercise this capability and even then only up to their allocation. All agents are able to revise their behaviour (*rev_behaviour*) according to perceived (sanctioning) events. In the role of a *member*, an agent is furthermore empowered to *vote* for a resource allocation method, place a *request* for resources and *appeal* against (subjectively) unwarranted sanctions and it holds fluents to account for the number of *offences* and the level of sanctions (*sanction_level*). What actions an agent decides to take and when depends on the environment and is explained in Section 6.2 together with the algorithm in Table 1.

According to the level of sanctions that an agent holds, the *gatekeeper* is empowered to *exclude* the member from the institution or permitted to *eliminate* a non-member according to the exclusion method (*exMethod*). The *monitor* monitors the appropriation of members (and non-members) and reports (*report* and *report_out*, respectively) any detected violation to the head. The *head* is in charge of calling for a vote (*cfv*), obliged to *declare* the resource allocation method and able to *allocate* resources to requesting members of the institution. In case of reported violations the head is empowered to *sanction* an agent, but also to *uphold* a sanction if it had been appealed against.

Finally, there is a class for the *institution* itself, with brute facts *Bf* like the *resource_level* and a corresponding *refill* function, or the impact of unintentional violations (*unintent_violation*). The institutional facts *If* include the chosen resource allocation method (*raMethod*), the access control and exclusion methods (*acMethod*, *exMethod*), the method to determine the winner of a vote (*wdMethod*) and the alternative dispute resolution method (*adrMethod*). A specific choice of these methods is part of one specification instance. The parameters of interest that are modified during the experiments are the 6 individual principles in combination with the frequency with which monitoring is taking place inside and outside the boundaries (*monitoring_freq*, *monitoring_freq_out*).

6.1.2. States. As the agents are considered heterogeneous, their individual characteristics (in this case the $compliancy\_degree$) lead to different behaviour when reacting to changes in the environment ϵ and specification instances l.

The four possible states are $active$ ***member***, $inactive$ ***member***, $active$ ***non-member*** and $inactive$ ***non-member***, see Figure 6. At the start of each run the $acMethod$ defines which of the agents become active members of the institution, the others become active non-members (Principle 1). Depending on the principles in use and how many resources an agent appropriates in relation to its allowances (Principle 2), the activity status of that agent can be changed.

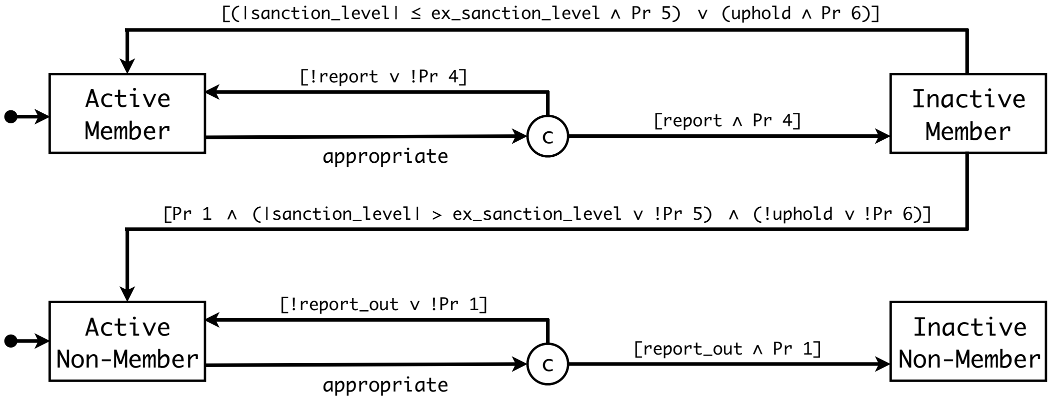

Fig. 6: Agent statechart

A member remains active, if it does not appropriate more resources than it is allowed to or if its misbehaviour has not been detected by the monitor and thus not reported ($!report$). Any detected non-compliance is reported (Principle 4) and the agent's status is set to inactive. The head is only permitted to set the agent's status to active again in the following two cases: The agent had not been sanctioned more often than the $ex\_sanction\_level$ (Principle 5 is being used) and has served its sentence (different penalties may apply for different sanction levels), or the sanction has been upheld due to a successful appeals procedure (Principle 6). In all other cases, the agent is excluded and assigned the active non-member status. That agent is not empowered to become a member again at any subsequent point in time.

An active non-member is not permitted to appropriate resources at all. Should it nevertheless be reported doing so (Principle 1), its status is set to inactive non-member, i.e. the agent is eliminated (for experimental purposes).

6.1.3. Algorithm. The classes in Figure 5 describe all the actions that agents are empowered to perform when occupying a particular role. These actions are determined by a change of the environment $\epsilon = \langle Bf, If \rangle$, i.e. how agents perceive it, and what is allowed or required according to the protocol. The sequence of the testbed's possible actions and events is described in Algorithm 1.

Initially, we set Principle 1 to active (true) and assign the roles of several *members* A and *non-members* A', a *head* C, *gatekeeper* G and *monitor* B, i.e. we $include$ all the agents at start. The time t is set to 0 and the resource pool is filled to the maximum level $P = P_{max}$.

ALGORITHM 1: Algorithm for the CPR testbed.

```
Principle1 ← $true$ ;                                          #existence of the institution
$\forall a \in A$.initially $role\_of\ (a, I) = member$ ;
$\forall a \in A'$.initially $role\_of\ (a, I) = non\text{-}member$ ;
$\exists c \in A$.initially $role\_of\ (c, I) = head$ ;
$\exists g \in A$.initially $role\_of\ (g, I) = gatekeeper$ ;
$\exists b \in A$.initially $role\_of\ (b, I) = monitor$ ;
$t \leftarrow 0$ ;  $P \leftarrow P_{max}$ ;                   #full resources
repeat
    if Principle3 then
        $cfv\ (c, raMethod, I)$ ;                              #call for votes
        $\forall a \in A.vote\ (a, X, raMethod, I)$ ;          #vote for allocation method
        $declare\ (c, W, raMethod, I)$ ;                       #allocation method is $W$
    else
        $declare\ (EXT, W, raMethod, I)$ ;                     #periodical external declaration
    endif
    if Principle2 then
        $\forall a \in A.demand\ (a, R', I)$ ;                 #resource request
        $\forall a \in A.allocate\ (c, a, R, I)$ ;             #resource allocation
    endif
    $appropriate\ (A, R'', I)$ ;                               #resource appropriation
    $appropriate\ (A', R'', I)$ ;
    $P \leftarrow P - \sum_{A \cup A'} R''$ ;
    if Principle4 then
        $report\ (b, A, R'', I)$ ;                             #monitoring
        $P \leftarrow P - P_{mon}$
        if Principle5 then
            $sanction\ (c, A, S, I)$ ;                         #sanction at level $S$
        endif
        if Principle6 then
            $appeal\ (a, S, I)$ ;                              #appeals procedure
            $uphold\ (c, a, S, I) \vee reject\ (c, a, S, I)$ ; #dispute resolution
        endif
    endif
    $exclude\ (b, A, I)$ ;                                     #exclusion
    $report\_out\ (b, A', R'', I)$ ;                           #boundary monitoring
    $P \leftarrow P - P_{mon\_out}$ ;
    $eliminate\ (b, A', I)$ ;                                  #elimination
    $t \leftarrow t + 1$ ;
    $P \leftarrow \min(P_{max},\ P + P_{rep})$ ;               #replenishment
until $(P < 0) \vee (A == \emptyset) \vee (t > t_{max})$;
```

The algorithm then cycles over t until a maximum amount of time steps t_{max} is reached, until the resource is depleted ($P < 0$) or there are no members left in the institution ($\nexists A \in I$).

In each time step, all agents perform the following actions, depending on whether the corresponding principle is selected. If collective-choice arrangements are being used (Principle 3), the head calls for a vote (*cfv*) on the resource allocation method, the agents vote for their preferred method and the head declares the winner to be the new *raMethod*. Otherwise, if Principle 3 is not used, the method for the current cycle is declared by an external partner EXT who re-evaluates it periodically. In order to then make the agents appropriate in accordance with the *raMethod*, Principle 2 has to be selected. Active members are empowered to place their demands R' and the head will allocate them $R \leqslant R'$, such that the sum of demands does not exceed the current resource level P. Afterwards, all agents perform an $appropriate$ action, where even the non-members are (physically) able to appropriate from the common pool.

If Principle 4 is selected, the monitor will monitor a proportion of the current members and report any offences to the head. The head then applies a sanction to the agent, here the sanction level S increases by 1 with each offence. In case Principle 5 is selected, the sanctioning is graduated, so that the agent is only excluded when reaching a certain limit of sanctions ($ex\_sanction\_level$). When an agent's status is set to active after the duration of its sanction, the agent can revise its behaviour and choose to be more compliant in the future. In case Principle 6 is used, the sanctioned (i.e. inactive) member is empowered to appeal against the sanction and the head might uphold it. Note that Principle 4 concerns both monitoring *and* enforcement, and if neither of the Principles 5 or 6 is used, the agent is excluded for a first reported offence directly by the gatekeeper. Moreover, the non-members are being monitored and reported ($report\_out$) in a similar way, but always eliminated for a first offence. Every elimination will act as a warning to other non-members and makes them revise their behaviour as well.

At the end of a cycle, t increases and the resource is replenished with P_{rep}, but only up to P_{max}.

6.2. Experimental Parameters

In this section we describe how the parameters of the algorithm are set to run the experiments. The range of the parameters, e.g. the refill rates, have been chosen such that they avoid both super-abundance, whereby restrictions on the resource allocation are redundant, and prolonged insufficiency, whereby the system is in a state of permanent crisis. Instead, the parameterisation sees to it that periods of high replenishment ensure that there should be sufficient resources in the long term, provided the appropriators avoid depleting the common-pool during periods of low resource replenishment.

The refill rate is part of the brute facts of the environment, as well as the resource level and environmental factors that motivate unintentional appropriations. The resource *level* can be modified by the agents through appropriation actions, but the *replenishment* and other environmental factors are beyond the agents' control. The highest level the resource can reach is $P_{max} = 10\,000$ and the refill rates change every 50 time steps, varying between low ($\leqslant 0.35P_{max}$), moderate ($\leqslant 0.45P_{max}$) and high ($> 0.45P_{max}$). If unintentional errors are taking place, 5% of the members appropriate more or fewer resources than they intend to. Moreover, 50% of the members (or non-members) can be set to have a low $compliancy\_degree$, meaning they are able to appropriate up to 20% more resources than they were allocated (or appropriate without allocation).

Principle 1 is a requirement for the *existence* of an institution (no members, no institution) and is set to active at start. The changeable parameter of this principle is an either high or low level of non-member monitoring ($monitoring\_freq\_out = 10\%$ or 1%) at a cost of 5 resource units per observation. In this testbed, the amount of resource an agent is appropriating per timeslice is known to the monitoring agent who in turn only queries this information with the defined frequency. It is furthermore assumed, the cost of monitoring is met from the resources.

The procedure of role assignment is not the focus of this testbed, therefore we start with 120 agents and 100 of them are included as members into the institution, the remaining ones are assigned the role of a non-member. An occupant for each of the roles of head, gatekeeper and monitor are assumed to be present at all times and not specifically appointed.

Together with Principle 1, Principles 2 and 3 serve the purpose of *managing* the institution. To vote on a resource allocation method, an agent makes a trade-off between the current resource level P and its own 'greed' and votes for its preferred $raMethod$. This will be *queue* if the resource in that time step is regarded as plentiful (typically $\geqslant 0.75P_{max}$), otherwise it will be *ration*. The head uses *plurality* as $wdMethod$ to declare the outcome of the vote. In case the method is determined externally, EXT declares a new $raMethod$ every 50 time steps following the same process.

If $raMethod$ is *queue*, an agent demands $R'_t = 50$ on average, placed in 90% of the time steps, and joins a queue. The agent leaves the queue after having been allocated R'_t, which can happen in a subsequent time step only. If $raMethod$ is *ration*, an agent will be allocated $\min{(R_t, R'_t)}$, where R_t is: P_t divided by the number of agents that placed a demand > 0 in t. The agent makes a rough estimate of R_t and adapts its demand R'_t.

Principles 4–6 *protect* the institution from malicious behaviour: 10% or 1% of the member agents, depending on a high or low $monitoring\_freq$, are monitored (at a cost of 50 resource units per observation), and when reported, a sanction is applied. For a first offence the agent is excluded from demanding resources (its activity status is set to inactive) for five time steps, for a second offence for ten, then 15. For $sanction\_level > 3$, an agent is excluded from the institution (active non-member). If the agent decides to appeal against the sanction, then the head upholds it if the agent has not been reported the 30 previous time steps and the sanction is withdrawn.

The institution's two main goals are the maximisation of its lifespan t (the amount of time until either $P < 0$ or $\nexists A \in I$) and a trade-off between resource level and sufficient membership. A sufficient membership is important to protect the resource against outsiders (not to be outnumbered) and will prove useful if there is a cost of ownership to cover.

6.3. Experimental Results and Evaluation

The system specifications, as described in the last subsection, have been implemented in a bespoke simulator using C++. The EC axioms have been implemented as C++ functions, because the testbed is not concerned with computing the normative states, only physical actions can be performed if not permitted (i.e. mis-appropriation), and the narrative of event is processed in temporal order of happening. In total, we conducted 4 sets of experiments (see Figures 7–10) to show under what environmental circumstances which principles are required, and with what parameter settings.

Each of the experiments was performed over 100 trials. All the following graphs are shown over time, and the value of the resource level per active institution (left) as well as the number of active agents per active institutions (right) were averaged over all trials. A third curve (axis on the far right) shows the number of institutions out of 100 trials that were still active, i.e. their resource had not yet been depleted at a particular time step and there were members remaining in the institution. In each of the graphs, the refill rate (high=h, moderate=m, low=l) during a certain period of time is displayed in rectangles and the arrows of partial legends point towards the number of active institutions.

6.3.1. Existence and management principles. Figure 7 shows an institution in whose environment the behaviour of the member agents is compliant, no unintentional violation occurs, but 50% of the agents outside the institution appropriate from the resource illicitly. Under these almost 'ideal' circumstances, Principles 1, 2 and 3 are tested, whose main purpose is to manage the resource allocation.

Principle 1 is used in every trial to guarantee the existence of the institution. The members of the institution are clearly defined and the boundaries are enforced by monitoring non-members and sanctioning them for a detected appropriation. To test

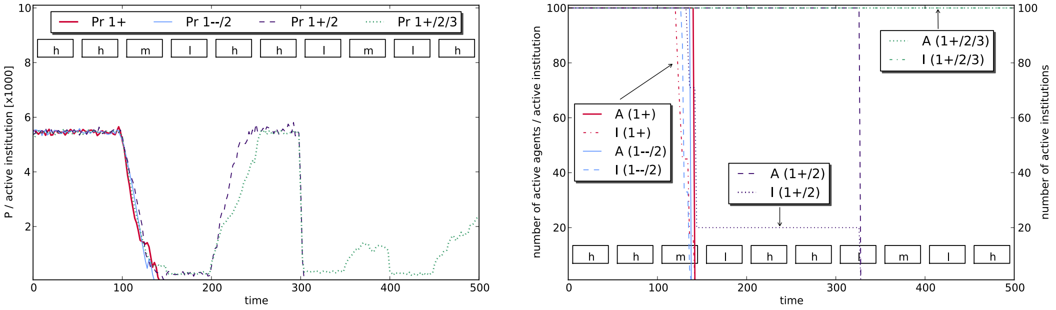

Fig. 7: Principles 1, 2 and 3: Lifespan of an ‘ideal’ institution I with all members compliant and no unintentional violation, 50% of the non-members are non-compliant.

the impact of enforcing the boundaries, this principle is implemented with two different settings of the boundaries’ monitoring frequency, high $(1+)$ and low $(1-)$. The first two runs shows the lifespan of I using Principle $1+$ only and the second with Principle $1-$ and 2. In both runs $(= 100$ trials each) all resources are depleted after around 140 time steps, when the refill rate lowers to medium. So neither a proper monitoring of the border nor the use of a resource allocation method (other than ‘first come, first served’) alone are enough for a sustainable resource and an enduring institution. If Principles $1+$ and 2 are used, the lifespan increases for a fifth of the trials by more than two times. Effectively, the institutions are now able to keep track of how many agents they have to allocate resources to and in what way. If Principle 3 is added, all institutions reach t_{max} without depleting the resource. Enabling the agents to choose the $raMethod$ according to their needs proved better than an external entity, that is unable to quickly respond to changes in the environment. This suggests that these three principles are indeed all needed to manage the resource appropriately.

6.3.2. Protection principles. Figure 8 shows four runs with Principles 1–3 active, then 1–4, 1–5 and finally Principles 1–6. From now on, Principle 1 is implemented with a high monitoring frequency.

This time, the environment of the institutions is different. 50% of the member agents do not comply to the appropriation rules and, moreover, unintentional violations occur. Also 50% of the non-members do not comply, as before.

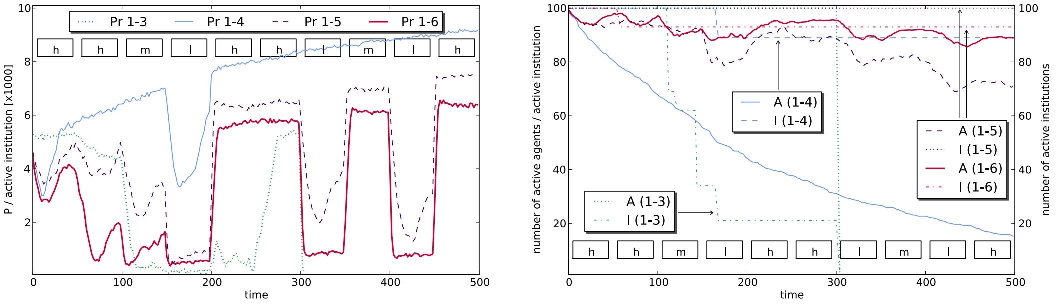

Fig. 8: Principles 4, 5 and 6: Lifespan of an institution I with 50% non-compliant members and non-members, and unintentional violation.

When only Principles 1–3 are selected, $4/5$ of the institutions' resources are depleted before $t = 140$ and the remaining ones survive little more than 300 time steps, due to the non-compliant behaviour of the members. Selecting Principle 4 in addition is supposed to identify exactly those members. As graduated sanctions are not in place, an agent is excluded from the institution for a first offence. This time, only very few institutions have their resource depleted when the refill rate becomes low, see the right side of Figure 8, but unfortunately far more than 50 agents are excluded from I due to unintentional violation. After the first low refill phase is passed, P reaches a high level as there are not many members left to (illicitly) appropriate from the resource.

When Principle 5 is added, no depletions occur. Non-compliant agents get the chance to revise their behaviour which results in many more members staying in the institution, although they are often sanctioned unwarrantedly due to unintentional violation. With Principle 6 in addition, this situation improves. Agents are now empowered to appeal and fewer agents are sanctioned or excluded. Very few of the institutions deplete the resource at an early stage. This is due to the fact that the appeals procedure is not errorless, and misjudgement of the nature of a member's violation (an agent may appeal against a sanction although it violated the rules intentionally) affects the member's propensity to revise its behaviour. Overall, for this experiment even more agents remain in the institution.

The results show that, when all six principles are in place, the institution is capable of dealing with different types of violations, intentional and unintentional, to ensure sustainability.

Furthermore, they align with Ostrom's and Hess' observation of suboptimal performance in the short, but better performance in the long run [Ostrom and Hess 2006]. This can be seen during the first 100 time steps in the runs with Principles 1–5 and 1–6, where the agents invest many of their resources into protection and therefore have a lower P, but prove more successful in the long run.

6.3.3. Setting the parameters. The next two sets of experiments were conducted to show that the implementation of the principles has to be carefully tuned to the environment ϵ of the institution, using Principle 4 as an example.

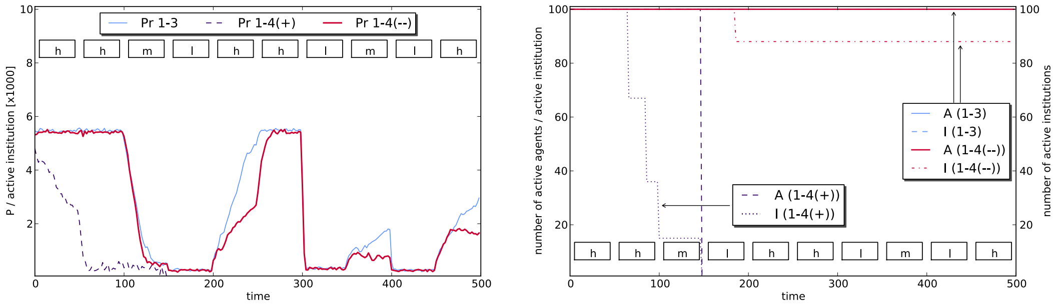

Fig. 9: Principle 4, high/low $(+/-)$ monitoring frequency: Lifespan of an institution I with 50% non-compliant non-members, all members compliant and no unintentional violation.

In Figure 9, three runs of the testbed are shown, where there were 50% non-compliant non-members, but all members are fully compliant and no unintentional violation occurs. As before in Figure 7, here in Figure 9 the institutions with Principles 1–3 selected do not deplete a single resource, $P \geqslant 0$ always. If we add Principle 4 using a high monitoring frequency, the related costs cause a depletion of resources in all trials before $t = 150$. If, instead, Principle 4 is implemented with a low monitoring frequency, only very few institutions deplete their resource completely, and the resource level of the remaining institutions is sufficiently high.

In contrast, Figure 10 shows three runs where, in addition, 50% of the members are non-compliant. This time, for the run with Principles 1–3 selected, all institutions end before $t = 160$ as the non-complying members are quickly depleting the resource. If we add Principle 4 with a low monitoring frequency, too few malicious agents are caught and the ones remaining in the institution cause a depletion of resource the sooner (almost 95% before $t = 175$) or later (the rest before $t = 410$). This time, expending additional resources to pay for the high monitoring frequency pays off. Only 30% of the institutions deplete their resource before $t = 175$ and no more afterwards, and about twice as many non-compliant agents are excluded from the institution.

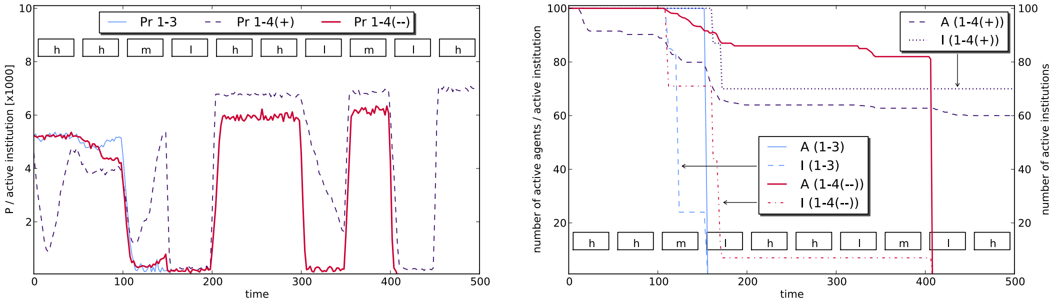

Fig. 10: Principle 4, high/low $(+/-)$ monitoring frequency: Lifespan of an institution I with 50% non-compliant members and non-members, no unintentional violation.

6.3.4. Evaluation. Table IV shows the benefit of each of the principles with respect to the agents' behaviour and environmental changes. Principles 1, 4, 5 and 6 are used to respond to unintentional and intentional violation of institutional facts If and to change the state of the agents according to Figure 6. Principles 2 and 3 allow the agents to (quickly) adapt to changes of the brute facts Bf in the environment, such as a shortage in resources.

Table IV: Comparison of the six principles and their benefits.

Principle	**Benefit**
1	robustness to intentional violation by outsiders
2	robustness to environmental variation
3	robustness to environmental variation
4	robustness to non-compliant behaviour
5	mitigation of intentional violation tolerance to unintentional violation
6	repair of unintentional violation

Depending on the behaviour of the agent population, i.e. when appropriation is performed by the rules, Principles 1–3 that manage the institution are sufficient for the institution to endure. As we relax the assumptions on the compliancy of the agents, Principles 1–3 are no longer enough and Principles 4–6 become necessary. Proposition **p3**, see Section 1, is supported by the results of the first two sets of experiments: all six principles together ensure an enduring institution with high membership as well as the sustainability of the resource.

Regarding the efficiency, we use the concept of *utilitarian social welfare (USW)* according to [Chevaleyre et al. 2007], and we measure the degree of efficiency by the sum of all individual utilities over the whole lifespan of the institution. In this context, the welfare corresponds to the amount of resource that all agents were appropriating together and was not spent on, for example, monitoring. The theoretical maximum is the amount of resource that could be appropriated at maximum over 500 time steps, taking the average demand and refill rates into account.

Table V: Efficiency of the individual experiments.

Experiments	**Environment**	**Principles**	**USW in 1000s**
Figure 7	100% member compliance + 50% non-member compliance + no unintentional violation	1+	596.79
		1-/2	601.77
		1+/2	780.19
		1+/2/3	1992.86
Figure 8	50% member compliance + 50% non-member compliance + unintentional violation	1-3	764.50
		1-4	1000.08
		1-5	1537.24
		1-6	1581.34
Figure 9	100% member compliance + 50% non-member compliance + no unintentional violation	1-3	1986.69
		1-4(+)	405.48
		1-4(-)	1829.26
Figure 10	50% member compliance + 50% non-member compliance + no unintentional violation	1-3	606.45
		1-4(+)	1306.13
		1-4(-)	734.33
theor. maximum			2001.89

In Table V, the utilitarian social welfare is shown for all of the previous experiments. As more principles are selected in accordance with the environment, the efficiency rises significantly. See, for example, the runs from Figure 7 where all members are compliant, i.e. Principles 4–6 are not needed, and the runs from Figure 8, where there are non-compliant members. For the next two sets of experiments, we can see that the most appropriate choice of principles in line with the local environment is the one with the highest efficiency, i.e. Principles 1–3 (no monitoring) in case when there are no non-compliant members, but Principles 1–4(+) (high level of monitoring) for the case with 50% non-compliant members. Where the institution does not need to spend resources on monitoring due to the agents' behaviour (Principles 1–3 in Figures 7 and 9), the efficiency almost reaches the theoretical maximum.

6.3.5. Limitations. One limitation of this testbed is that the EC specification has been used indirectly as a formal specification rather than directly as an executable specification. Given the number of events generated in each round, either a more efficient implementation of the Event Calculus is required (e.g. [Artikis et al. 2012]), or a testbed with an interface to an efficient rule engine, such as Drools[1] should be developed.

---

[1] www.jboss.org/drools/

Other limitations become evident in the last two sets of experiments, where a high or low monitoring frequency resulted in a better or worse outcome depending on the agents' behaviour. All principles have to be carefully implemented with respect to the prevailing environment, meaning that a trade-off with respect to a 'one size fits all' parameter cannot be made. We therefore need additional mechanisms for the institution to first learn what its environmental states are, and second respond to them appropriately. In addition, there are a wide range of independent variables, especially population distribution and resource level variation, all of which need to be analysed: this is a common challenge facing experimental evaluations of this type.

In Section 4.4 we mention a distance function d that is defined on the set of specification instances of $\mathcal{L}$. For example, this distance can be interpreted as a transaction cost, which then means that not only the running costs have to be taken into account when changing to a different $l \in \mathcal{L}$, but also the costs involved in the process of getting there. For instance, it might be expensive to call for a vote to change the monitoring frequency, but a lower frequency (and lower running cost) might pay off in the long run, if there are only very few malicious agents. These cost considerations have not been made in the current testbed.

The challenge is then to estimate the overall utility of changing l, as we do not know how long a certain instance is going to be suitable for the environment. To this end, we need mechanisms for runtime self-analysis; perceiving agents that reflect on events and deduct consequences from prospective actions; and agents that are self-aware and choose their policies in line with the environment, so to beneficially influence the behaviour of non-compliant agents.

7. RELATED WORK AND FURTHER RESEARCH CHALLENGES

Traditionally, the role of a software engineer has been to apply some methodology to implement a 'closed' system which satisfies a set of functional and non-functional requirements. Our problem is to engineer 'open' systems where the primary non-functional requirement, that the system should endure, is an *emergent* property, and is a side-effect of the *interaction* of components rather than being the goal of any of those components. More generally, unplanned emergent behaviour exhibited by complex socio-technical systems cannot easily be handled by top-down design methods. Thus the approach proposed here has much in common with other new design methods, for example *design for emergence* [Ulieru 2007], for systems that adapt and evolve, and where the design method specifically targets 'self-*' properties. Note that this is related to, but is a different target from, the emergence of norms, which has been studied, in the context of the the tragedy of the commons, from the perspective of socio-psychology [Schindler 2012], immergence [Conte et al. 2009], and metanorms [Mahmoud et al. 2011].

Our aim has been to leverage Ostrom's work for *agent-based software engineering*, but there is also related research from the perspective of *agent-based modelling*. This reveals many additional parameters to consider in developing experiments to test the emergent property of endurance. For example, Janssen *et al.* [2008] investigate whether or not people are prepared to invest their own resources in endogenous rule change, e.g. from open access to private property. We will also have to design experiments which consider the 'cost' of rule changes, the costs of monitoring and dispute resolution, and the impact this has on 'endurance'. In addition, Ostrom's original analysis has been extended to introduce more than 30 factors which influence endurance [Agrawal 2001], and we may need to enrich our model with these additional parameters. The meta-analysis of [Cox et al. 2010] also revealed a sub-division of three principles into two parts, and for some applications we may need to refine the axiomatisation to reflect these distinctions.

The results of this work contribute to the theory and practice of *electronic institutions*, as a paradigm for structured interactions using conventions or norms in an open environment. Several approaches have been proposed for the specification of such systems. For example, Esteva and colleagues [Esteva et al. 2002; Esteva et al. 2004] have devised a formal language (ISLANDER) to specify open MAS as electronic institutions and a middleware (AMELI) for subsequently executing such specifications. This work has also been used for agents to learn norms for adaptive organisations [Miralles et al. 2010]. Work from the field of computational organisation theory includes the work of Fox and colleagues on enterprise modeling, e.g. [Fox et al. 1998]. In this approach, a multi-agent organisation is viewed as a set of agents playing roles in which they are acting to achieve specific goals, according to various constraints defining the 'rules of the game'. The rules are formalised with the use of a dialect of the Situation Calculus [Pinto and Reiter 1993]. An interesting challenge for these approaches is to formalise the concept of institutionalised power, which is critical to the formalisation of Ostrom's principles in a computational setting, in the respective specification languages and see if the corresponding execution reproduces the experimental results of enduring institutions reported here.

There remain, however, many further research challenges, including the effect of Principle 8, the effect of different types of resource, full use of the specification space and the distance metric d, full use of the EC axiomatisation as an executable specification, issues of self-awareness and fairness, and practical applications. We briefly discuss these challenges here.

Although many related works on institutional action and institutionalised power parameterise their formal accounts with respect to an institution, for 'simplicity' or 'expediency' it is assumed that there is just one institution. However, the key feature of Principle 8 is that there are layered or encapsulated CPRs, or multiple CPRs operating in the same space. This is why we have included the parameter I in the fluents and actions of our EC specification, as a placeholder for further work on systems of systems of CPRs. We plan to implement an 'asynchronous' version of the testbed based on Ostrom's notion of a *decision arena* and a more efficient EC dialect. All resource allocation decisions would take place in one decision arena, all dispute resolution procedures would take place in another, and so on. This would allow the application of operational-choice rules, collective-choice rules, etc. to overlap and interleave, rather than all being resolved within one timeslice.

There are a number of experiments on the notion of 'nesting' that need can be performed in a setting with Principle 8 active. We want to investigate nesting in four directions. Firstly, there is the full nesting of operational-choice within social collective-choice, including the role assignment of the *head*, and the selection of the $wdMethod$, for example, and the formalisation of the decision arenas at each nested level. The second direction is the embedding of institutions within larger institutions, rather than the single layer model implemented here, to form the system of systems identified by Ostrom. The third direction involves third parties and other dependencies which can lead to other, more complex, supply chains. For example, a CPR system for water can, in times of drought, create a dependency for a CPR for food distribution, and so on. Finally, the fourth direction involves the formalisation of constitutional-choice rules, which we propose to model as constraints on the specification space, i.e. different constitutions will give different valid and invalid specification instances, different cost functions, and different constraints on the DoF values.

The current work has focussed on exogenous resource supply. Axtell outlines a dynamic model for team formation based on evolutionary game theory [Axtell 2002], in which a set of agents attempt to form a stable coalition. This work analysed the conditions under which agents cooperate, and demonstrated that groups become unstable beyond a certain size due to free riding. This argument was extended in [Axtell 2007] to volatile populations of agents, who leave and join teams based on a local view of utility rather than through a cyclical up/down time. There is scope to include such behaviour in our system. This work also argues that the mathematical complexity of such volatile systems precludes any analytic results. An axiomatisation in the Event Calculus supports off-line tasks like proving properties, and supports direct computational implementation for experimental investigation, when the randomness in the system makes the system behaviour inherently unpredictable.

In [Pitt et al. 2011c], we have also considered the case of endogenous resource provision and cluster formation using the Linear Public Good game [Gaechter 2006]. This game is concerned with provision and appropriation of resources, in which individual utility is maximised by free-riding (i.e. providing nothing and appropriating an equal share) if everyone else provides a full share, but minimised if everyone tries to free-ride. The experimental results in this case showed that the distributed self-organisation according to Ostrom's principles was robust even to initially non-compliant populations (i.e. agents which were biased towards free-riding), and that its behaviour approximated the theoretically ideal centralised solution with 'perfect' agents which always complied.

In addition to exogenous and endogenous resources, there are also concerns about how the principles of self-organisation work in relation to information resources, and multiple, re-usable and/or indivisible resources. Understanding information resources such as digital libraries or wikipedias as a *knowledge commons* has been proposed [Ostrom and Hess 2006]. The use of auctions [Murillo et al. 2011], traditionally considered as a centralised solution, to provide a mutable resource allocation method for an open decentralised system requiring provision and appropriation of re-usable, multiple and/or indivisible resources, is being investigated.

There are also aspects of Artikis' framework [Artikis 2011] which are currently under-utilised. This includes the use of a topological space to express the 'distance' between two specification instances and its relationship to the 'cost model' of institutional change defined by Ostrom, and their joint impact on the role assignment protocol (i.e. some agents may be 'trusted' more than others to occupy a role).

This also introduces additional scope for partial observation and unintentional error, and examining in more detail how the running costs of monitoring and dispute resolution affect compliance. A representation of cost using the distance function d could also be used to learn to correlate the specification space with the trajectory of environment change to determine the preferred specification instance for a particular environmental state (in a sequence of states). However, this also needs to take into consideration the incentives and costs of persuading other agents to agree to change the specification [Poteete et al. 2010].

Note that various ways to define this distance between specification instances may be used. The choice of distance function is application-specific. For example, in some cases it is possible to have a total order of specification instances while in other applications only a partial order is possible. Depending on the requirements of the application under consideration, the distance function may be based on a total order or a partial order of specification instances (see, for example, [Horling and Lesser 2008; Jilani et al. 2001]).

Looking further ahead to applications, we are interested in applying Ostrom's principles to service-oriented computing and SmartGrids. In service-oriented computing, increasing attention is being paid to applications of cloud computing for enterprise management and business delivery, in particular the real-time on-demand provisioning of Software-as-a-Service (SaaS), Infrastructure-as-a-Service (IaaS), and so on. It is interesting to cast this problem as a non-cooperative game of group formation and resource management problem [Ardagna et al. 2011]. In that work, a game theoretic approach is proposed, based on competing SaaS providers managing IaaS provider capacity. It would be worthwhile to investigate the effects of the clusters themselves as competing entities, each offering flat-rate, on-demand and spot-market resource access, and to model this from an institutional perspective. This will entail a more fine-grained representation of service level agreements, electronic contracts, and quality of service; and a more refined model of resource allocation based on flat rates, on-demand and spot market service provision.

Much of the work on policies for cloud computing focuses on the issue of security and defining security policies [Birman et al. 2009; Carminati et al. 2009]. Since role assignment is important for some aspects of security (i.e. access control), we believe this work complements that research. Furthermore, in real-world cloud computing environments there may be multiple access control policies and even multiple gatekeepers. The non-monotonic reasoning about institutional facts supported by the EC supports straightforward modelling of such situations. Moreover, alternative dispute resolution methods can be formalised in the same language and used to implement conflict-resolution protocols.

On sustainability, we draw attention to the MAELIA project [Boulet et al. 2009], which is building a multi-agent platform to model the interaction, from a network perspective, of agents, actions and norms on renewable resources. Their emphasis is understanding how to analyse and optimise policy with respect to sustainability, and their representation of norms is not grounded in an action language. Scaling up the system described here and deploying it for demand-side management of physical resources (where the resource consumers are also resource providers) is a substantial and significant challenge for further application of these ideas to management of autonomic power systems and SmartGrids (cf. [Strbac 2008]).

The final challenge concerns lifting self-organising institutions to self-aware institutions. The experiments showed that the implementation of Ostrom's principles has to be carefully designed and modified to be congruent with the current state of the environment, but also the nature of the population. For example, there was no 'efficiency' gain from expensive and/or extensive monitoring of a compliant population, and it was easy to deplete a resource by unnecessary over-monitoring, see Figure 9. Therefore, there has to be some form of institutional awareness of the nature of the population and the purpose of rules in achieving a macro-level goal (i.e. not depleting the resource). Only a self-aware institution can be sure to avoid path dependencies and ensure continual successful self-organisation that not only produces efficient but also 'fair' outcomes. However, fairness itself is subjective to a particular context and there are several metrics the agents can choose from [Lan et al. 2010; de Jong and Tuyls 2011].

8. SUMMARY AND CONCLUSIONS

In summary, we applied a methodology for sociologically-inspired computing to a (pre-formal) theory of socio-economics. First, we presented the problem of resource allocation in open, embedded systems in terms of Ostrom's socio-economic theory of self-governing institutions for common-pool management (**p1**). We then cast this theory into a formal model and a computational framework for dynamic specification of , and gave a complete axiomatisation in the Event Calculus of six of the eight principles proposed by Ostrom for enduring institutions (**p2**). This provided a deeper conceptual understanding of the protocols and institutionalised powers underlying the principles, and the inter-connection between them. On this basis, we designed and implemented an experimental testbed to investigate the dynamic behaviour of the system. The results showed that Ostrom's principles are necessary and sufficient conditions to create self-organising electronic institutions that endure and manage common-pool resources sustainably (**p3**).

In conclusion, we have shown that an institution-based approach to the problem of dynamic resource allocation in an open, embedded and resource-constrained system is feasible, and has particular advantages when long-term endurance of the distribution mechanism is more important than short-term 'optimality'. Our experiments have shown that a collaborative process of institutional change can achieve improved performance through self-organisation, in what might otherwise be considered as an N-player non-cooperative game. While there are many domain-specific issues to address, there are significant opportunities to apply this model to cloud computing, or to SmartGrids and the idea of virtual power plants for the allocation of distributed energy resources [Pudjianto et al. 2008]. We also believe that the work reported here has laid the foundations to address further challenges, for example in the automation of decision-support systems for adaptive institutions to address climate change [RCEP 2010], institutions for smarter infrastructure management, including water and transport as well as energy, and a deeper investigation into the development of institutions that are not only self-organising, but are also *self-aware*.

ACKNOWLEDGMENTS

This research was partially supported by EU FP7 project AWARE Grant Number 257154 and UKIERI project SA08-037. This paper is an extended version of [Pitt et al. 2011b] and we are grateful to the anonymous reviewers of both that earlier version and this extended version for their very helpful comments. We are also grateful to Elinor Ostrom, Claudia Brink and Michael Cox of the Workshop in Political Theory and Policy Analysis at Indiana University for further discussion on this research.

REFERENCES

AGRAWAL, A. 2001. Common property institutions and sustainable governance of resources. *World Development 29,* 10, 1649–1672.

ANDREWS, P., POLACK, F., SAMPSON, A., STEPNEY, S., AND TIMMIS, J. 2010. The cosmos process, version 0.1: A process for the modelling and simulation of complex systems. Technical Report YCS-2010-453, University of York.

ARDAGNA, D., PANICUCCI, B., AND PASSACANTANDO, M. 2011. A game theoretic formulation of the service provisioning problem in cloud systems. In *WWW '11*. ACM, New York, 177–186.

ARTIKIS, A. 2011. Dynamic specification of open agent systems. *Journal of Logic and Computation*, doi: 10.1093/logcom/exr018.

ARTIKIS, A. AND SERGOT, M. 2010. Executable specification of open multi-agent systems. *Logic Journal of the IGPL 18,* 1, 31–65.

ARTIKIS, A., SKARLATIDIS, A., PORTET, F., AND PALIOURAS, G. 2012. Logic-based event recognition. *Knowledge Engineering Review*.

AXTELL, R. 2002. Non-cooperative dynamics of multi-agent teams. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems*, C. Castelfranchi and W. Johnson, Eds. ACM, New York, 1082–1089.

AXTELL, R. 2007. What economic agents do: How cognition and interaction lead to emergence and complexity. *The Review of Austrian Economics 20,* 2, 105–122.

BIRMAN, K., CHOCKLER, G., AND VAN RENESSE, R. 2009. Toward a cloud computing research agenda. *SIGACT News 40,* 2, 68–80.

BOULET, R., MAZZEGA, P., AND JOUVE, B. 2009. Environmental, social and normative networks in the maelia platform. In *Proc. ICAIL Worshop Legal & Decision Support Systems*, M. Poblet, U. Schild, and J. Zeleznikow, Eds. 83–93.

BRAMS, S. AND TAYLOR, A. 1996. *Fair Division – From cake-cutting to dispute resolution*. CUP, Cambridge.

CARMINATI, B., FERRARI, E., AND PEREGO, A. 2009. Enforcing access control in web-based social networks. *ACM Trans. Inf. Syst. Secur. 13,* 1, 1–38.

CHEVALEYRE, Y., ENDRISS, U., LANG, J., AND MAUDET, N. 2007. A short introduction to computational social choice. In *Proceedings of the 33rd Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM)*. Number 4362 in LNCS. Springer-Verlag, Berlin, Heidelberg, 51–69.

CONTE, R., ANDRIGHETTO, G., AND CAMPENNÍ, M. 2009. The immergence of norms in agent worlds. In *Engineering Societies in the Agents World X*, H. Aldewereld, V. Dignum, and G. Picard, Eds. LNCS Series, vol. 5881. Springer, 1–14.

COX, M., ARNOLD, G., AND VILLAMAYOR TOMÁS, S. 2010. A review of design principles for community-based natural resource management. *Ecology and Society 15,* 4, 38.

DE JONG, S. AND TUYLS, K. 2011. Human-inspired computational fairness. *Autonomous Agents and Multi-Agent Systems 22,* 1, 103–126.

DOMINGOS, P. AND LOWD, D. 2009. *Markov Logic: An Interface Layer for Artificial Intelligence*. Morgan & Claypool Publishers.

DOUSSON, C. 1996. Alarm driven supervision for telecommunication networks – on-line chronicle recognition. *Annales des Télécommunication 51,* 9-10, 501–508.

EDMONDS, B., GILBERT, N., GUSTAFSON, S., HALES, D., AND KRASNOGOR, N., Eds. 2005. *Socially Inspired Computing. Proceedings of the Joint Symposium on Socially Inspired Computing*. AISB.

ESTEVA, M., DE LA CRUZ, D., AND SIERRA, C. 2002. Islander: an electronic institutions editor. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems*. 1045–1052.

ESTEVA, M., RODRÍGUEZ-AGUILAR, J. A., ROSELL, B., AND ARCOS, J. 2004. Ameli: An agent-based middleware for electronic institutions. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multi-agent Systems*. 236–243.

FARRELL, A., SERGOT, M., SALLÉ, M., AND BARTOLINI, C. 2005. Using the event calculus for tracking the normative state of contracts. *International Journal of Cooperative Information Systems 14,* 2–3, 99–129.

FIROZABADI, B. S. AND SERGOT, M. 2004. Contractual access control. In *Proc. 10th Int. Workshop on Security Protocols (2002)*, B. Christianson, B. Crispo, J. Malcolm, and M. Roe, Eds. Number 2845 in LNCS. Springer-Verlag, Berlin, Heidelberg, 96–103.

FOX, M., BARBUCEANU, M., GRUNINGER, M., AND LIN, J. 1998. An organisation ontology for enterprise modeling. In *In Simulating Organizations: Computational Models of Institutions and Groups*, M. Prietula, K. Carley, and L. Gasser, Eds. Menlo Park CA: AAAI/MIT Press, 131–152.

GAECHTER, S. 2006. Conditional cooperation: Behavioral regularities from the lab and the field and their policy implications. Discussion Papers 2006-03, The Centre for Decision Research and Experimental Economics, School of Economics, University of Nottingham.

HEGSELMANN, R. AND KRAUSE, U. 2002. Opinion dynamics and bounded confidence: Models, analysis and simulation. *Artificial Societies and Social Simulation 5,* 3.

HORLING, B. AND LESSER, V. R. 2008. Using quantitative models to search for appropriate organizational designs. *Autonomous Agents and Multi-Agent Systems 16,* 2, 95–149.

JANSSEN, M., GOLDSTONE, R., MENCZER, F., AND OSTROM, E. 2008. Effect of rule choice in dynamic interactive spatial commons. *International Journal of the Commons 2,* 2, 288–312.

JILANI, L. L., DESHARNAIS, J., AND MILI, A. 2001. Defining and applying measures of distance between specifications. *IEEE Transactions on Software Engineering 27,* 8, 673–703.

JONES, A. AND SERGOT, M. 1996. A formal characterisation of institutionalised power. *Journal of the IGPL 4,* 3, 427–443.

KATSH, E. 2006. Online dispute resolution: Some implications for the emergence of law in cyberspace. *Lex Electronica 10,* 3.

KOWALSKI, R. AND SERGOT, M. 1986. A logic-based calculus of events. *New Generation Computing 4*, 67–95.

KREMER, I. AND NYBORG, K. 2004. Divisible-good auctions: The role of allocation rules. *RAND Journal of Economics 35,* 1, 147–159.

LAN, T., KAO, D., CHIANG, M., AND SABHARWAL, A. 2010. An axiomatic theory of fairness in network resource allocation. In *INFOCOM*. 1343–1351.

MAHMOUD, S., KEPPENS, J., LUCK, M., AND GRIFFITHS, N. 2011. Norm establishment via metanorms in network topologies. In *IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology,*. 25–28.

MIRALLES, J., LÓPEZ-SÁNCHEZ, M., ESTEVA, M., AND MORALES, J. 2010. A simulator for organisation-centred mas adaptation in p2p sharing networks. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems*. 1615–1616.

MURILLO, J., BUSQUETS, D., DALMAU, J., LÓPEZ, B., MUNOZ, V., AND RODRÍGUEZ-RODA, I. 2011. Improving urban wastewater management through an auction-based management of discharges. *Environmental Modelling & Software 26*, 689–696.

OSTROM, E. 1990. *Governing the Commons*. CUP, Cambridge.

OSTROM, E. AND HESS, C. 2006. A framework for analyzing the knowledge commons. In *Understanding Knowledge as a Commons: From Theory to Practice*, C. Hess and E. Ostrom, Eds. MIT Press, Cambridge, MA, 41–82.

PINTO, J. AND REITER, R. 1993. Temporal reasoning in logic programming: A case for the situation calculus. In *Proceedings of the Tenth International Conference on Logic Programming*. 203–221.

PITT, J., KAMARA, L., SERGOT, M., AND ARTIKIS, A. 2006. Voting in multi-agent systems. *Computer Journal 49,* 2, 156–170.

PITT, J., RAMIREZ-CANO, D., DRAIEF, M., AND ARTIKIS, A. 2011a. Interleaving multi-agent systems and social networks for organized adaptation. *Computational and Mathematical Organization Theory 17,* 4, 344–378.

PITT, J., SCHAUMEIER, J., AND ARTIKIS, A. 2011b. The axiomatisation of socio-economic principles for self-organising systems. In *Fifth IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO)*. 138–147.

PITT, J., SCHAUMEIER, J., AND ARTIKIS, A. 2011c. Coordination, conventions and the self-organisation of sustainable institutions. In *Agents in Principle, Agents in Practice*. LNCS Series, vol. 7047. Springer-Verlag, Berlin, Heidelberg, 202–217.

POTEETE, A. R., JANSSEN, M. A., AND OSTROM, E. 2010. *Working Together: Collective Action, the Commons, and Multiple Methods in Practice*. Princeton University Press.

PUDJIANTO, D., RAMSAY, C., AND STRBAC, G. 2008. Microgrids and virtual power plants: concepts to support the integration of distributed energy resources. *Proc IMechE, A: J. of Power and Energy 222*, 731–741.

RAYA, M. AND HUBAUX, J.-P. 2007. Securing vehicular ad hoc networks. *Journal of Computer Security 15,* 1, 39–68.

RCEP. 2010. 28th report: Adapting institutions to climate change. Royal Commission on Environmental Protection (Chairman: John Lawton), The Stationery Office Limited, UK.

ROBERT, S. C., ROBERT, H., EVANS, W. J., HONEMANN, D. H., AND J., B. T. 2000. *Robert's Rules of Order, Newly Revised, 10th edition*. Perseus Publishing, Cambridge, MA.

SANDHU, R., FERRAIOLO, D., AND KUHN, R. 2000. The NIST model for role-based access control: Towards a unified standard. In *5th ACM Workshop Role-Based Access Control*. RBAC '00. ACM, New York, 47–63.

SCHINDLER, J. 2012. Rethinking the tragedy of the commons: The integration of socio-psychological dispositions. *Journal of Artificial Societies and Social Simulation 15,* 1, 4.

STEELS, L. AND BROOKS, R. 1994. *The artificial life route to artificial intelligence: Building Situated Embodied Agents*. New Haven: Lawrence Erlbaum Ass.

STRBAC, G. 2008. Demand side management: Benefits and challenges. *Energy Policy 36,* 12, 4419–4426.

TIDEMAN, N. 2006. *Collective decisions and voting: the potential for public choice*. Ashgate Publishing Ltd, Aldershot, UK.

ULIERU, M. 2007. Evolving the 'DNA blueprint' of eNetwork middleware to control resilient and efficient cyber-physical ecosystems. In *2nd International Conference BIONETICS*. 41–47.

Received ; revised ; accepted