Coordination, Conventions and the Self-organisation of Sustainable Institutions

Jeremy Pitt[1], Julia Schaumeier[1], and Alexander Artikis[1,2]

[1] Department of Electrical & Electronic Engineering, Imperial College London, SW7 2BT, UK
[2] National Centre for Scientific Research "Demokritos" Athens 15310, Greece

Abstract. Applications where autonomous and heterogeneous agents form opportunistic alliances, which require them to share collective resources to achieve individual objectives, are increasingly common. We model such applications in terms of self-governing institutions for shared resource management. Socio-economic principles for enduring institutions are formalised in a logical framework for dynamic specification of norm-governed systems. The framework is implemented in an experimental testbed to investigate the interplay of coordination in a social dilemma with mutable conventions of an institution. Experimental results show that the presence of conventions enables the norm-governed system to approximate the performance of a theoretically ideal system. We conclude that this approach to self-organisation can provide the foundations for implementing sustainable electronic institutions.

1 Introduction

Applications in which autonomous and heterogeneous agents form opportunistic alliances, which require them to share collective resources in order to achieve individual objectives, are increasingly common. Examples include vehicular networks [14], service-oriented systems such as cloud computing [1], and demand-side infrastructure management for water [6], energy [16], and so on. These examples are all open, distributed and resource-constrained. However, we are unable to 'privatise' the system, otherwise it would no longer be open, nor to '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 [12]. By 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 so can be adapted to suit the environment in which the system is embedded. This might itself be changed by exogenous events.

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

D. Kinny et al. (Eds.): PRIMA 2011, LNAI 7047, pp. 202–217, 2011.

should not only suit the environment but also result in a 'fair' outcome. Thirdly, even a fair distribution has to be sustainable in the long term; in other words, the rules also have to ensure that the institution itself is somehow enduring.

In the investigation of these criteria, this paper is organised as follows. Section 2 reviews the background to this work. Using a methodology for engineering socio-technical systems [9], Sections 3 and 4 develop a formal characterisation of self-governing institutions as dynamic norm-governed systems. In Section 5 describes experiments to evaluate the interplay of coordination in an iterated n-player game with mutable conventions of an institution. Results show that using conventions enables the institution to approximate the performance of a theoretically ideal system. Related and further work is discussed in Section 6, and we conclude in Section 7 that this approach to self-organisation provides the foundations for implementing electronic institutions whose properties of compliance pervasion, fairness and endurance support sustainability for CPR management.

2 Background

This section reviews the background to the current work, including the work on CPR management of Ostrom [12], institutionalised power [10], the dynamic specification of norm-governed systems [2], and the linear public good game [7].

2.1 Self-governing Institutions

Ostrom [12] observed that common pool resource (CPR) management problems have often been resolved in human societies 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" [12, 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 [12, 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. In 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 a 'fixed' set of rules at a higher level. Secondly, lower level rules were easier and less 'costly' to change than the higher level rules, thus increasing the stability of strategies and expectations of those individuals having to interact with others in the context of the institutional setting.

Ostrom also observed that there were occasions when the institutions were *enduring*, and others where they were not. Accordingly, eight principles of institutions were identified for *self*-management of common pool resources (CPR) to endure. Of these, three were:

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. For example, an appropriate rule that allows everyone an unrestricted claim on resources is not congruent with the environmental condition where those same resources are scarce.
3. Collective choice arrangements: in particular, those affected by the operational rules participate in the selection and modification of those rules.

It is necessary to identify who is a *member* of the institution, and who is not, as it is precisely the members of the institution who are those affected by modification of the rules. We also need to distinguish specific members who are *empowered* to enact, announce and enforce these modifications.

2.2 Institutionalised Power and Roles

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 agent that occupies a designated role, that is *empowered* to make those decisions.

Therefore, we need to represent the concepts of role, role assignment [15], and *institutionalised power* [10]. 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 norm-governed specifications.

2.3 Dynamic Specifications

Artikis [2] defined a framework that allowed agents to modify the rules or protocols of a norm-governed system at runtime. This framework defined 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.

A specification of a norm-governed system can be (partially) given by defining the permissions, prohibitions and obligations of the agents in the system, and the sanctions and enforcement policies that deal with the performance of prohibited actions and non-compliance with obligations [3].

The protocol stack allowed agents to modify the rules or protocols of a norm-governed system at runtime. This framework defined a set of object level protocols, and assumed 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, we need to define who is, and who is not, a *member* of an institution, where agents can join an institution if they satisfy certain criteria, and can be excluded if they do not comply to the rules. We specify two types of method, one for access control and another for exclusion. The type of access control method is *acMethod*, which can be *attribute-based*, whereby if the applicant satisfies certain qualification criteria then it is automatically admitted, or *discretionary*, i.e. an applicant must satisfy another agents'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 a agent.

Each type of method is a DoF, and with two values for each method, this gives four possible *specification instances*. This the basis for defining a *specification space* $\langle T, d \rangle$, where T is the set of all possible specification instances and d is a function which defines a 'distance' between any pair of elements in T.

2.4 Linear Public Good (LPG) Game

CPR management by an institution requires that each agent provides to and appropriates resources from the common pool. The agents must comply with the rules concerning 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 in a CPR is considered as a linear public good (LPG) game [7]. This problem 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 $i, i \in \{1, \ldots, n\}$ decides independently to contribute resources $r_i \in [0, 1]$ to the public good. The contributions from the whole cluster are summed and the payoff u_i for each player i is given by:

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

The first term represents the payoff from the public good (the 'public payoff'), 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 a and b represent the relative value of the public/private payoffs respectively. If the conditions on a and b hold, a rational but selfish agent has the incentive to contribute 0 to the public good, i.e. free riding, 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 Formal Characterisation

In this section, we describe a methodology for sociologically-inspired computing [9], apply it to cast Ostrom's definition of an institution (Section 2.1) as a dynamic norm-governed specification (Section 2.3), and derive a formal model of a multi-agent system to play the n-player iterated linear public good game.

3.1 Methodology

A methodology for sociologically-inspired computing is illustrated in Figure 1.

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 [12] comes into this category, as it is an evidence-based theory of enduring institutions but without formalism. 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

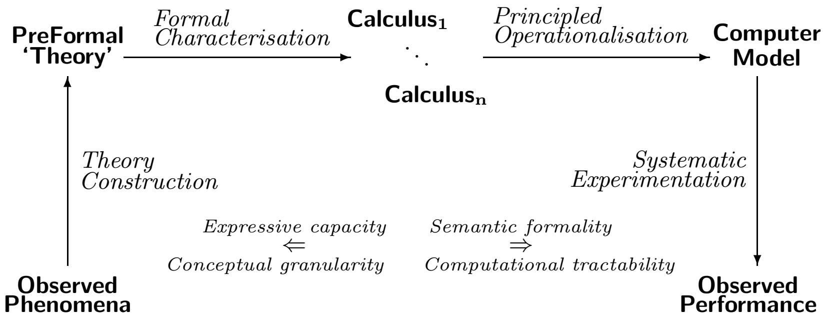

Fig. 1. Sociologically-Inspired Computing

role of the calculus: expressive capacity or conceptual granularity with regard to 'theory'; computational tractability or semantics with regard to implementation. The step of *principled operationalisation* embeds such formal representations in simulations which include detailed implementation of individual agents.

3.2 Institutions as Dynamic Specifications

The three elements of Artikis' framework [2] were a norm-governed specification, a protocol stack, and a specification space.

Firstly, the institutional rules of Ostrom are characterised as a norm-governed specification. As such, the specification will define the following 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) [11] is used as the calculus for formal characterisation.

Secondly, the nesting of operational-choice rules within 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 we will specify in the 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.

The Artikis framework originally defined the specification space as a metric space. In practice, we find this too restrictive and instead of a metric space we represent the set of specification instances T as nodes on a graph with a constant 'distance' k between any two nodes, i.e. $\forall l_1, l_2 \in T, d(l_1, l_2) = k$. The

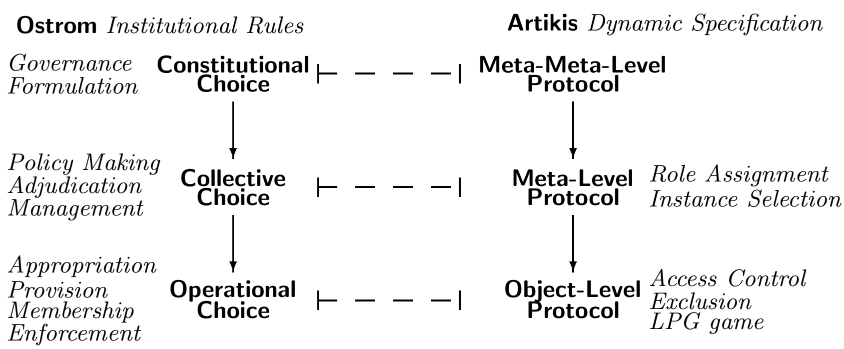

Fig. 2. Institutional rules as Protocol Stack

specification space used here is given below in Section 4.4. (Note that we will not use d further in these experiments; however, we find it convenient to retain it for further work in representing the 'cost' of modifying operational, collective and constitutional choice rules.)

3.3 Formal Model

We will now instantiate a formal model of an institution for the LPG game. Let $\mathcal{IC}_t$ be a multi-agent system at time t defined by:

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

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 norm-governed system specification (defining a specification space T);
- G is the LPG game;
- d is a distance function defined on specification instances of T.

Each institutional cluster $I_t \in \mathcal{I}_t$ is given 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 T; and
- ϵ is the cluster's local environment, a pair $\langle Bf, If \rangle$ with Bf the set of 'brute' facts whose values are determined by the physical state, including the average contribution made by members to the cluster; 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 intuitive idea is that at each time-point t, the agents in $\mathcal{A}$ will form into clusters $\mathcal{I}$ using the access control method. Each cluster plays a linear public good game, where the members either comply or defect by contributing more or less resources than the cluster average. After the game, non-compliance may be punished by exclusion according to the operational rules.

A specific type of institutional fact recorded in ϵ is which agent is empowered to perform a certain role in each cluster. We identify four roles: *member*, which is the standard role for membership of a cluster in order to participate in G; *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.

Therefore the set $\mathcal{L}$ contains the following two rules for role assignment, with (in parenthesis) the role responsible for its enactment and enforcement:

$$(\mathit{gatekeeper}) \;\; \mathit{ocr}_1 : \mathcal{M}^c \times \mathit{acMethod} \rightarrow \mathit{Bool}$$
$$(\mathit{monitor}) \;\; \mathit{ocr}_2 : \mathcal{M} \times V(\cdot)_{a \in I} \times \mathit{exMethod} \rightarrow \mathit{Bool}$$

where $V(\cdot)_{a \in I}$ is a set of expressed preferences on an issue by each *member* agent in cluster I, where $I \in \mathcal{I}$.

The operational choice rule ocr_1 is applied by the *gatekeeper* to map an application to join from an agent not in $\mathcal{M}$ (i.e. the set complement $\mathcal{M}^c$) to a boolean outcome depending on the access control method. A true result means the applicant can be assigned the role of *member*. Similarly, the rule ocr_2 is applied by the *monitor* to map an agent in $\mathcal{M}$ that did not comply with the rules of the LPG game to a boolean outcome using the exclusion method.

4 Action Language Specification

In this section, we illustrate the axiomatisation of the rules in $\mathcal{L}$ using the Event Calculus (EC) [11], and define the graph for the 'specification space'. A summary of the EC is given in [2], and a full EC specification of six of Ostrom's principles in [13]. For space reasons, we do not review EC or reproduce those axioms.

4.1 Fluents (Institutional Facts)

Some of the institutional facts, represented as fluents F of the EC, are as follows. The multi-valued fluent *role_of* has the value *head*, *gatekeeper*, *monitor* or *member* if the agent occupies the associated role, and so participates in the LPG game G for a cluster I, and has the value *none* otherwise.

The multi-valued fluent *acMethod* determines which access control method the gatekeeper must use in determining *member* role assignment. Its value is either *attribute* or *discretionary*. The multi-valued fluent *exMethod* determines which exclusion method the monitor must use in determining *member* exclusion (note that the the jury method requires a winner-determination method. Technically this is a mutable DoF but we will assume it is fixed at plurality here). Three other fluents record the (institutionalised) powers, permissions and obligations of each agent. These are all institutional facts in *If*. The real-valued fluent *cluster_average*, a physical fact in *Bf*, records the average contribution of the agents to the public good.

In this specification, we stipulate that agents can occupy only one role in a cluster and that agents can be members of only one cluster. It is straightforward to modify the specification (and the testbed used for the experiments, described in the next section), so that agents can occupy more than one role and be members of more than one cluster, but neither choice fundamentally affects the issue being investigated, i.e. self-regulating sustainable institutions.

4.2 Member Role Assignment (ocr_1)

The dynamic specification of the operational choice rule ocr_1 is given by a role-assignment protocol for membership. An agent can apply for membership to a cluster I if it does not occupy a role in any other cluster:

$$\begin{array}{l} apply(A, I) \quad \mathsf{initiates} \quad applied(A, I) = true \quad \mathsf{at} \quad T \quad \leftarrow \\ \qquad \mathbf{not} \quad role\_of(A, \_) = member \quad \mathsf{holdsAt} \quad T \quad [\text{etc.}] \end{array}$$

The gatekeeper agent is empowered to admit the agent, to the cluster, by an *assign* action, depending on the access control method.

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

If the *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 *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 the external conditions.

4.3 Member Exclusion (ocr_2)

The *monitor* is empowered to exclude a *member* that does not comply with the rules of the game G. For each iteration of G, agents should contribute resources in the interval $[ave_I, 1]$ to comply, where ave_I is the average contribution of resources from the previous iteration (the value of the fluent $cluster\_average(I)$). For each iteration, an agent's default provision is 0, and if it does not *provide* an average (or greater) provision then it is sanctioned:

$$\begin{array}{l}
provide(A, R, I) \ \ \mathsf{initiates} \ \ provision(A, I) = R \ \ \mathsf{at} \ \ T \ \ \leftarrow \\
\quad \mathbf{pow}(A, provide(A, R, I)) = true \ \ \mathsf{holdsAt} \ \ T \\
provide(A, R, I) \ \ \mathsf{initiates} \ \ sanctioned(A, T, I) = true \ \ \mathsf{at} \ \ T \ \ \leftarrow \\
\quad \mathbf{pow}(A, provide(A, R, I)) = true \ \ \mathsf{holdsAt} \ \ T \ \ \wedge \\
\quad provision(A, I) = R \ \ \mathsf{holdsAt} \ \ T \ \ \wedge \\
\quad cluster_average(I) = Ave \ \ \mathsf{holdsAt} \ \ T \ \ \wedge \ \ R < Ave \\
\mathbf{pow}(A, provide(A, R, I)) = true \ \ \mathsf{holdsAt} \ \ T \ \ \leftarrow \\
\quad role_of(A, I) = member \ \ \mathsf{holdsAt} \ \ T
\end{array}$$

If the *exMethod* is discretionary, then the *monitor* agent G can exclude the applicant A (or not) as it decides. If the *exMethod* is jury, then the *monitor* must have called for a vote on the issue of the exclusion of A:

$$
\begin{array}{l}
\mathit{exclude}(G, A, \mathit{member}, I) \quad \mathsf{initiates} \quad \mathit{role\_of}(A, I) = \mathit{none} \quad \mathsf{at} \quad T \quad \leftarrow \\
\qquad \mathbf{pow}(G, \mathit{exclude}(G, A, \mathit{member}, I)) = \mathit{true} \quad \mathsf{holdsAt} \quad T \\
\mathbf{pow}(G, \mathit{exclude}(G, A, \mathit{member}, I)) = \mathit{true} \quad \mathsf{holdsAt} \quad T \quad \leftarrow \\
\qquad \mathit{role\_of}(G, I) = \mathit{monitor} \quad \mathsf{holdsAt} \quad T \quad \wedge \\
\qquad \mathit{exMethod}(I) = \mathit{discretionary} \quad \mathsf{holdsAt} \quad T \\
\mathbf{pow}(G, \mathit{exclude}(G, A, \mathit{member}, I)) = \mathit{true} \quad \mathsf{holdsAt} \quad T \quad \leftarrow \\
\qquad \mathit{role\_of}(G, I) = \mathit{monitor} \quad \mathsf{holdsAt} \quad T \quad \wedge \\
\qquad \mathit{exMethod}(I) = (\mathit{jury}, \mathit{WDM}) \quad \mathsf{holdsAt} \quad T \quad \wedge \\
\qquad \mathit{ballot}(\mathit{exclude}(A), I) = V \quad \mathsf{holdsAt} \quad T \quad \wedge \\
\qquad \mathit{winner\_determination}(\mathit{WDM}, V, \mathit{true}) \\
\mathbf{per}(G, \mathit{exclude}(G, A, \mathit{member}, I)) = \mathit{true} \quad \mathsf{holdsAt} \quad T \quad \leftarrow \\
\qquad \mathit{role\_of}(G, I) = \mathit{monitor} \quad \mathsf{holdsAt} \quad T \quad \wedge \\
\qquad \mathit{sanctioned}(A, T', I) = \mathit{true} \quad \mathsf{holdsAt} \quad T \quad \wedge \quad T' < T
\end{array}
$$

Note that the *monitor* is empowered to exclude any *member*, but it is only *permitted* to exercise that power when that member has been sanctioned (and, when, the exclusion method is *jury*, only when the vote is in favour of exclusion). This means that when the monitor excludes an agent, that agent really is excluded and it has no role in the institution. However, an excluded agent can appeal against an invalid use of the power, the *monitor* could be removed from the role, and so on. This is a higher-order effect which is beyond the scope of the current paper. Furthermore, voting and winner determination has been studied in this context in [13], but is not considered further here.

4.4 Specification Space

These rules in $\mathcal{L}$ effectively define four DoF (degrees of freedom): the selection of the *acMethod* and the selection of the *exMethod*, and the assignment to the *gatekeeper* role and the assignment to the *monitor* role. Meta-level protocols for role assignment and instance selection can be specified in the EC, as above, but for space constraints are omitted here.

Since no agent can occupy both roles *monitor* and *gatekeeper* in a cluster I with n members, it follows that there are $4n^2-4n$ possible specification instances. Rather than dynamically computing the entire space for each cluster and trying to determine the 'optimal' configuration, we separate the 'specification instance' selection function into two dimensions.

For the first dimension, the decision of which agent to assign to the role of *monitor* or *gatekeeper*, we define a family of preference functions, some based on relevant properties of the agent (e.g. compliance probability, time already spent in the role, etc.) and some not (e.g. random, nominative proximity, etc.). Each agent is associated with a subset of these functions, and applies them when

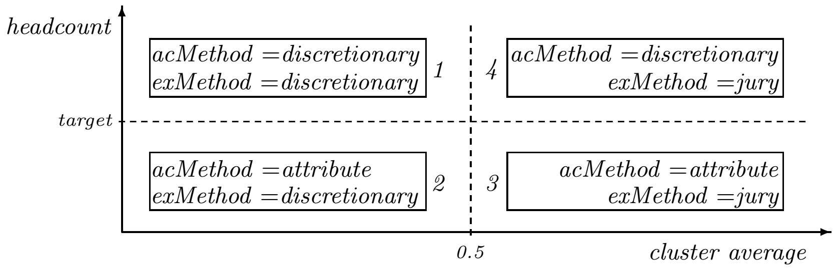

Fig. 3. Specification Space

voting for either *monitor* or *gatekeeper*. The *head* is empowered to assign the role to the agent with the most votes according to a winner determination method.

For the second dimension, the selection of *acMethod* and *exMethod*, we define two criteria. The first criteria is a *target* membership: this value is a trade-off between total cost of ownership (which is too high if the *headcount* is less than the target) and the quality of service (which it too low if the *headcount* is more than the target). The second criteria is the average probability of compliance in the LPG game. For each agent, we define a probability distribution for voting for a change in the specification according to these criteria. As a result, the selected specification instance falls into one of the quadrants *1–4* as shown in Figure 3.

5 Experimental Results

This section describes the implementation of a testbed and experimental results, evaluating the performance of an institutional approach to the LPG game.

5.1 Testbed Implementation

The control loop for the testbed is shown in Algorithm 1. A run starts with the random generation of a population of M agents and N clusters, and sets the time-point t to 0. One agent is assigned to the *head* role in each cluster.

To introduce an element of volatility, each agent (except the *head*) is associated with a random cycle of length x time-points, of which it is 'present' for y (time-points) and 'absent' for z, such that $x = y + z$. In step 5, those agents that transition from absent to present are given the status *present*, but have no role in any cluster; those that transition from present to absent leave their cluster and whatever roles they occupy, and their status is *absent*.

In step 6 and 7, the member agents vote for a monitor or gatekeeper, if one or both of those roles have been vacated. Then, *present* agents which are not members of a cluster (either because they became present from step 5, they failed to get into a cluster or they were excluded from a cluster in the previous time-point) apply to a cluster. The *gatekeeper* applies the *acMethod* as presented in Section 4.2 and assigns the applicant to the role of *member* or it is rejected. In step 9, member agents play the LPG game within each cluster, and the cluster average is updated. Non-complying agents may then be excluded as per Section 4.3 using the operational *exMethod*. Afterwards, each cluster's *acMethod* and *exMethod* are updated by a vote of its members, the time-point is incremented and the cycle repeats.

Algorithm 1. Control Loop for CPR testbed.

```
generate_agents(M, A)
generate_clusters(N, C)                          %assign (designate) head
t ← 0
repeat
   update_present(A)
   gatekeeper_role_assignment(A, C)              %role assignment by vote
   monitor_role_assignment(A, C)                 %role assignment by vote
   member_role_assignment(A, C)                  %use acMethod, Section 4.2
   public_good_game(A, C)                        %play LPG game, Section 2.4
   member_exclusion(A, C)                        % %use exMethod, Section 4.3
   update_clusters(C)                            %update acMethod, exMethod by vote
   t ← t + 1
until t = ...
```

5.2 Agent Strategies

When a population of agents is generated, a bundle of information is associated with each agent. This includes its name, up-time, down-time, initial cluster and role assignment (may be none), and its strategy for the LPG game. This strategy is given by a probability of complying with the rules of the game. Therefore the contribution $r_{i,t}$ that an agent i makes at time t is given by:

$$\begin{aligned} r_{i,t} &= Ave_{(t-1)} + rnd(1) \cdot (1 - Ave_{(t-1)}), \quad \text{if} \quad rnd(1) \geqslant pc_{i,t} \\ &= rnd(1) \cdot Ave_{(t-1)}, \quad \text{otherwise} \end{aligned}$$

where $pc_{i,t}$ is the probability of i's compliance at t and $Ave_{(t-1)}$ the average cluster contribution from the last time-point. As a result the contribution is in the interval $[Ave, 1]$ if a random number generated in the interval $[0, 1]$ is greater than the probability of compliance, and is in the interval $[0, Ave]$ otherwise.

The agents update their probability of compliance in the next time-point according to a form of social influence. Letting $|\,I\,|$ denote the *headcount* for the number of agents in cluster I and $|\,I\,|^+$ denote the number of agents in I which complied in the current round, then:

$$\begin{aligned} pc_i(t+1) &= pc_i(t) + \alpha \cdot (1 - pc_i(t)), \quad \text{if} \quad (|\,I\,|^+ \,/\, |\,I\,|) \geqslant 0.5 \\ &= pc_i(t) - \beta \cdot pc_i(t), \quad \text{otherwise} \end{aligned}$$

where α and β are globally defined coefficients in $[0, 1]$ determining the rate of positive and negative *reinforcement* respectively. If a majority of agents complies in the current round, then the likelihood of each one complying in the next round is increased, and vice versa.

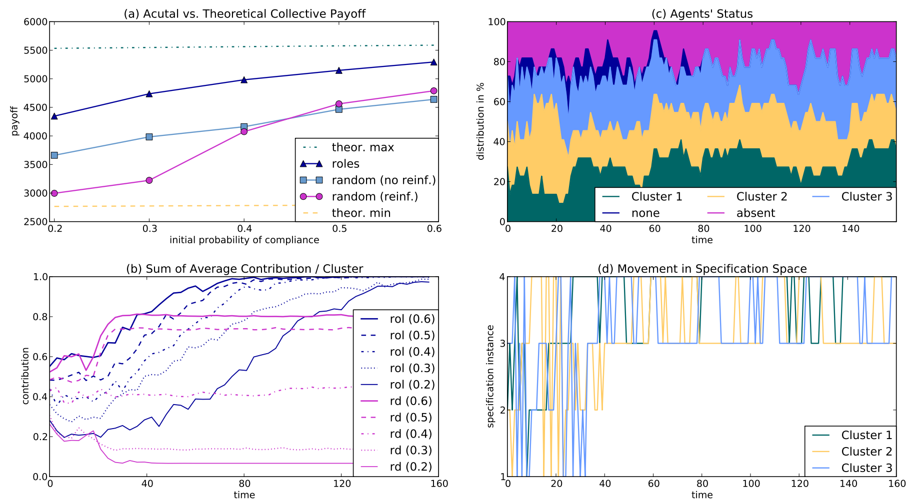

Fig. 4. Experimental Results

5.3 Evaluation

The experiments were run with 22 agents and 3 clusters, the same population distribution was used for 10 runs each of 160 time-points, and the results averaged. The independent variable was the initial probability of compliance for each agent, for the entire population. The dependent variables were the combined payoffs and the combined cluster averages. The compliance reinforcements rates were $\alpha = \beta = 0.05$ and the LPG game coefficients were $a = 2$ and $b = 1$.

Figure 4(a) shows the results for 5 populations starting with a probability of non-compliance 0.2, through to 0.6. There is a theoretical maximum collective reward at each time-point, which is the number of agents 'present', independent of whether they were in a cluster or not, multiplied by 2 (the coefficient a), and a theoretical minimum, i.e. the number of present agents (since the coefficient b is 1). Three sets of 10 runs are shown: for a population with role assignment and with updating the probability of compliance (reinforcement), without role assignment (i.e. random) but with reinforcement, and without either role assignment or reinforcement (i.e. random with a fixed probability of compliance). The graph shows that with roles and reinforcement, even at low levels of initial compliance, the cluster performs better than random, with or without reinforcement. As the initial compliance increases, the overall payoff increases and starts to approximate the payoff of the theoretical ideal. Reinforcement without rules is about as effective as without either reinforcement or rules at higher levels of initial compliance, but actually worse at lower levels of initial compliance.

Figure 4(b) displays the cumulative moving cluster averages for the same agent populations, totalled over the 10 runs each, for clusters with role assignment and reinforcement and for clusters without role assignment (random) but with reinforcement. The graph shows that, even at low levels of initial compliance, the agents using role assignment and reinforcement were able to form stable and sustainable clusters, with full compliance pervasion as each agent is contributing 1.0. Without role assignment, each run tends either to full non-compliance or full compliance, and as expected this happens around 50% of the time with initial compliance probability of 0.5/0.6, and much less often at lower levels of initial compliance. Thus the sum of the average cluster contributions, taken over 10 runs, is then less than full compliance.

Figure 4(c) shows a typical distribution of agents to clusters. Initial probability of non-compliance is 0.5, coefficients α, β, a, b as before. Each cluster has a target membership of 6, or approximately one-third of the expected total number of agents expected to be present at any one time-point. The total number of present but non-member agents (as a result of exclusions or rejections) reduces to zero by about time-point 100; after that, the distribution of present agents to clusters is more or less even.

Finally, Figure 4(d) shows the change of access control and exclusion methods in the specification space. About the time agents stop being excluded, the access control and exclusion methods oscillate between specification instances 3 and 4, i.e. the cluster average is high so the access control method sticks at discretionary, and the exclusion method varies according to the number of agents present. This confirms that the cluster average was so high it was only the headcount that was affecting which specification instance the agents were using.

The robustness of the institutions and the distribution of rewards suggests that the institution is, in some sense, both ‘fair’ and ‘enduring’. On that basis, we contend that the management of the shared contributed resources, which is the responsibility of the institutions, is sustainable.

6 Related and Further Work

The linear public good game used here was studied in [17], and took an evolutionary algorithms approach to solving the social dilemma it presented, in contrast to the explicit representation of rules and institutions in this work.

Axtell outlines a dynamic model for team formation based on evolutionary game theory [4], 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 [5] 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. Our 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 service-oriented computing, there is increasing attention 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 [1]. 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.

Our aim here has been to leverage Ostrom's work for *agent-based software engineering*, but there is 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, [8] investigates whether or not people are prepared to invest their own resources in endogenous rule change, e.g. from open access to private property. There is much scope for investigating richer agent strategies in this context.

On sustainability, we draw attention to the MAELIA project [6], 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 work (cf. [16]).

Finally, the testbed offers several directions for further research. Currently, the testbed implements the EC axioms and generates the actions, but does not put the narrative through an EC engine. We are investigating use of the cached EC for this purpose. This also introduces scope for partial observation and unintentional error, and examining how the costs of monitoring and dispute resolution affect compliance. A representation of cost could also be used to explore the specification space and strategies to use of the distance function d to determine a preferred specification instance. This also relates to Ostrom's comments about the cost of changing operational- and collective-choice rules.

7 Summary and Conclusions

In summary, we applied a methodology for sociologically-inspired computing to a (pre-formal) theory of socio-economics. Our aim was to cast self-governing institutions for common pool resource management in the framework of dynamic (norm-governed) specifications. The resulting formal model was given a complete axiomatisation in the Event Calculus and an experimental testbed was designed and implemented to investigate the dynamic behaviour of the system. The results showed that the distributed self-organising system was robust even to initially non-compliant populations, that its behaviour approximated the theoretically ideal centralised solution with 'perfect' agents, and that the distribution of rewards indicated that the institution was, in some sense, 'fair' and sustainable.

Our technical conclusion is that the application of institutional rules and institutional change for CPR management has a beneficial impact on autonomous and autonomic multi-agent systems. The challenge is to apply these ideas in socio-technical multi-agent systems as part of a sustainable infrastructure for common pool resource management, such as water and energy.

References

1. Ardagna, D., Panicucci, B., Passacantando, M.: A game theoretic formulation of the service provisioning problem in cloud systems. In: WWW 2011, pp. 177–186 (2011)
2. Artikis, A.: Dynamic protocols for open agent systems. In: Proc. AAMAS 2009, pp. 97–104. IFAAMAS (2009)
3. Artikis, A., Sergot, M., Pitt, J.: Specifying norm-governed computational societies. ACM Transactions on Computational Logic 10(1), 1–42 (2009)
4. Axtell, R.: Non-cooperative dynamics of multi-agent teams. In: Castelfranchi, C., Johnson, W. (eds.) Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems, pp. 1082–1089 (2002)
5. Axtell, R.: What economic agents do: How cognition and interaction lead to emergence and complexity. The Review of Austrian Economics 20(2), 105–122 (2007)
6. Boulet, R., Mazzega, P., Jouve, B.: Environmental, social and normative networks in the maelia platform. In: Poblet, M., Schild, U., Zeleznikow, J. (eds.) Proc. ICAIL Worshop Legal & Decision Support Systems, pp. 83–93 (2009)
7. Gaechter, S.: 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 (2006)
8. Janssen, M., Goldstone, R., Menczer, F., Ostrom, E.: Effect of rule choice in dynamic interactive spatial commons. International Journal of the Commons 2(2), 288–311 (2008)
9. Jones, A., Pitt, J., Artikis, A.: On the analysis and implementation of normative systems: Towards a methodology: In: Cranefield, S., Noriega, P. (eds.) PreProceedings COIN@AAMAS 2011, pp. 47–56 (2011)
10. Jones, A., Sergot, M.: A formal characterisation of institutionalised power. Journal of the IGPL 4(3), 427–443 (1996)
11. Kowalski, R., Sergot, M.: A logic-based calculus of events. New Generation Computing 4, 67–95 (1986)
12. Ostrom, E.: Governing the Commons. CUP (1990)
13. Pitt, J., Schaumeier, J., Artikis, A.: The axiomatisation of socio-economic principles for self-organising systems. In: Proceedings SASO 2011 (2011)
14. Raya, M., Hubaux, J.P.: Securing vehicular ad hoc networks. Journal of Computer Security 15(1), 39–68 (2007)
15. Sandhu, R., Ferraiolo, D., Kuhn, R.: The NIST model for role-based access control: Toward a unified standard. In: 5th ACM Workshop Role-Based Access Control, RBAC 2000, pp. 47–63 (2000)
16. Strbac, G.: Demand side management: Benefits and challenges. Energy Policy 36(12), 4419–4426 (2008)
17. Yamashita, T., Axtell, R., Kurumatani, K., Ohuchi, A.: Investigation of mutual choice metanorm in group dynamics for solving social dilemmas. In: Kurumatani, K., Chen, S.-H., Ohuchi, A. (eds.) IJCAI-WS 2003 and MAMUS 2003. LNCS (LNAI), vol. 3012, pp. 137–153. Springer, Heidelberg (2004)