Holonic Institutions for Multi-scale Polycentric Self-governance

Ada Diaconescu[1] and Jeremy Pitt[2](✉)

[1] Telecom ParisTech, 46 Rue Barrault,75634 Paris Cedex 13, France
[2] Imperial College London, Exhibition Road, London SW7 2BT, UK
j.pitt@imperial.ac.uk

Abstract. Effective institutions are key to the success of self-governing systems, yet specifying and maintaining them can be challenging, especially in large-scale, highly dynamic and competitive contexts. Political economist Elinor Ostrom has studied the conventional arrangements for sustainable natural resource management and derived from these eight design principles for self-governing institutions. One principle, *nested enterprises*, is straightforwardly expressed, but is arguably structural rather than functional, and so is more resistant to declarative specification; yet it also appears to be critical to the effectiveness of complex compositional systems. In this paper, we converge the ideas of holonic systems with electronic institutions, to propose a formalisation of this principle based on *holonic institutions*. We show how holonic institutions provide a structural framework for nested enterprises, which can be designed as composite systems of systems. This, we believe, is compatible with Ostrom's ideas for polycentric governance of complex systems. We use a case study in energy distribution to illustrate these ideas.

Keywords: Electronic institutions · Holonic architectures · Multi-agent systems · Self-organising systems · Polycentric governance · Smartgrids

1 Introduction

Based on extensive fieldwork examining successful, and unsuccessful, instances of common-pool resource management, Ostrom [17] identified eight common features of the successful instances, some of which were missing from the unsuccessful ones. She then posited these features as design principles for the supply (endowment) of self-governing institutions for sustainable resource management.

These principles are extensively documented [17] and only briefly reminded here (Sect. 4), with the exception of the eighth principle, concerning *nested enterprises*. This principle states that institutions, which consist of conventional rules, are nested within each other, with provision and appropriation systems operating locally at a small-scale (base level) and being organised into multiple layers at larger-scales over wider geographical regions (higher levels).

The principle itself is straightforwardly expressed and is arguably structural rather than functional – i.e., it is more concerned with the structural relationships between institutions than the purposeful functions those institutions are

A. Ghose et al. (Eds.): COIN 2014, LNAI 9372, pp. 19–35, 2015.
DOI: 10.1007/978-3-319-25420-3_2

intended to deliver. As such this principle has proven more resistant to declarative specification than the other principles [20]. Yet, it also appears to be critical to the effective functioning of complex compositional systems operating across multiple scales, with multiple objectives and intricate interdependencies.

For example, reducing global Carbon emissions could be considered as a collective action problem consisting of country-level actors, but regulation by the Kyoto protocol has failed to meet its targets. Indeed, Ostrom herself posed the question: are large-scale collective action problems, with correspondingly large-scale outcomes, better addressed by large-scale government policies [15]? For Ostrom, the answer was equivocal; but generally in the case of climate change, somehow, the system of nested enterprises is failing to provide the appropriate distribution of policy formation, decision-making and self-governance. Therefore, Ostrom argued, policies made at national and international level also required local and regional action and enforcement. Governance had to be *polycentric* – i.e. composed of multiple centres of decision-making [14] – enabling complex, multi-scale systems to cope with complex, multi-criteria problems.

There is, however, a fairly well-established understanding in utilising *holonic architectures* to address complex systems issues, such as scalability, heterogeneity and dynamic adaptability, via the recursive coordination of processes that operate at different granularity levels. Holonic architectures and their key role in creating viable complex systems were introduced by Simon [26], refined by Koestler [9], and progressively adopted in software systems engineering. For instance, Simon argues that holarchy "is one of the central structural schemes that the architect of complexity uses" [26]. Hence, the central question addressed in this paper is: *can holonic architectures be used to implement Ostrom's nested enterprises institutional design principle for polycentric governance?*

Accordingly, we converge the ideas of holonic systems with electronic institutions implementing executable forms of Ostrom's principles [20], and propose a formalisation of the nested enterprises design principle based on *holonic institutions*. The paper is structured as follows. Section 2 reviews the background and motivation for this work, while Sect. 3 introduces the convergence of holonic architectures and electronic institutions. This is the basis for a preliminary study of holonic institutions in Sect. 4, with an illustrative case study of community energy systems in Sect. 5. We conclude that this indicates how holonic institutions could provide a composite system of systems architecture for nested enterprises and inter-linked organisations which, we believe, is compatible with Ostrom's ideas for polycentric governance of complex systems.

2 Background and Motivation

2.1 Formalising Ostrom's Principles

The primary aim of using Ostrom's principles as the basis for electronic institutions was to address the problem of resource allocation in open computing systems and networks. In open systems, the components effectively form a common pool of resources (CPR) and specify conventional rules concerning provision to, and appropriation of, resources from the common pool. In the absence of a centralised component strictly enforcing the rules, and the possibility of sub-ideal behaviour (from accidental operation, to free-riding and intentional malice), Ostrom's design principles were proposed to supply self-governing institutions which supported sustainable resource management.

In the experiments of [20], six of Ostrom's eight design principles were specified in computational logic. It was shown that, as more principles were added, the electronic institutions moved along the spectrum from failure (usually depleted the resource) through fragility (sometimes depleted, sometimes sustained the resource) to sustainability (usually sustained the resource). This replicated the findings reported by Ostrom in [17, p. 180, Table 5.2].

Of the other two principles, the seventh concerned *no external authorities*, which was effectively implemented since there were no external authorities (although it was *not* shown, in [20], that some form of external authority disrupted an institution's capability to sustain a resource). The eighth principle, *nested enterprises*, was *NR* (not relevant, to borrow the classifier from Table 5.2 cited above): this principle only concerned "CPRs that are parts of larger systems" [17, p. 90]. In [20], there was only a single, base-level CPR.

2.2 The Eighth Principle: Nested Enterprises

The eighth principle is highly significant for multiple institutions, more complex systems, or electronic institutions for socio-technical systems. Here, Ostrom's fieldwork indicates a dependence between multiple CPRs. For example, in irrigation systems, there is a CPR for appropriation of water. Given water's tendency to flow downhill, the expectation would be that those at the 'top end' would appropriate all the water, leaving nothing for those at the 'bottom end'. However, this does not (always) happen: it turns out there is a second CPR, for maintenance of the irrigation system, which the top-enders cannot manage on their own. If they appropriate all the water, the bottom-enders don't provision to the maintenance CPR. Therefore it is successful collective action in one CPR which provides the *social capital* [16] for successful collective action in the other; so in fact there are two, asynchronous but co-dependent, CPRs whose inter-operation serves to sustain the resource.

Similarly, in SmartGrids for power management, there has been a shift from the traditional model of *predict and provide* to demand-side management – i.e., given the power available, schedule the demand to fit. This shift has partly been motivated by the increase in stochastic generators and the perceived impossibility of centralised scheduling of millions of dispatchable generators under such constraints. One solution is to form a hierarchy of autonomous virtual power plants (AVPP) [10], and to delegate scheduling to each AVPP in the hierarchy. However, these works mostly focus on the control functions necessary to achieve predefined goals – e.g., avoiding load peaks and maximising provider revenues; based on rules that are known in advance – e.g. switching on and off equipment such as heaters and fridges. They do *not* consider *how* the institutional rules that guide these controls are negotiated, specified or evolved, by members of the socio-technical system, for achieving justice, fairness or conformance objectives.

Finally, while these electronic institutions have been inspired by formalising observations about social systems, it is an open question what happens if such institutions are injected back into the social system, to form a socio-technical system. One example would be a socio-technical system for demand-side power management, or better, demand-side self-organisation. However, such a system would inevitably be part of a much grander socio-technical system, a system of nested enterprises, with base-level concerns (over price, stability and availability, say) to the user, and country-level concerns over Carbon emissions at the top. In other words, this is a system of multi-scale, multi-objective nested enterprises subject to possibly competing policy constraints.

2.3 An Example

Consider a single entity producing and consuming resources (a *prosumer*). On its own, it may strike a balance between production and consumption; alternatively at times it may generate more or less resources than are required, which may be wasteful or risk causing a blackout. To avoid these problems, the prosumer can coordinate with others and pool their resources, subject to the self-organisation and mutual agreement of the *rules of engagement*. These rules constitute an institution, in the sense of Ostrom [17].

Suppose that, as in [19], the institution operates in time slices, during which each agent generates resources, computes its resource requirements, provisions resources to the common pool, receives an allocation, and makes an appropriation. There are several *operational-choice* rules involved, for example concerning provision. There are (at least) two alternatives: firstly, that a prosumer in the institution should provision all the resources that it generates to the common pool; secondly, that it only needs to provision any excess beyond its own requirements.

In the framework in [19,20], these rules could be formalised in the Event Calculus (EC) [11] as shown below.

$$
\begin{aligned}
&\mathbf{obl}(H, \mathit{provide}(A, P_a, I)) = \mathit{true} \quad \mathsf{holdsAt} \quad T \quad \leftarrow \\
&\quad \mathit{role}(A, I) = \mathit{prosumer} \quad \mathsf{holdsAt} \quad T \quad \wedge \\
&\quad \mathit{rule}(I, \mathit{provision}) = \mathit{all} \quad \mathsf{holdsAt} \quad T \quad \wedge \\
&\quad \mathit{generated}(A) = P_a \quad \mathsf{holdsAt} \quad T \\
&\mathbf{obl}(H, \mathit{provide}(A, P_a, I)) = \mathit{true} \quad \mathsf{holdsAt} \quad T \quad \leftarrow \\
&\quad \mathit{role}(A, I) = \mathit{prosumer} \quad \mathsf{holdsAt} \quad T \quad \wedge \\
&\quad \mathit{rule}(I, \mathit{provision}) = \mathit{excess} \quad \mathsf{holdsAt} \quad T \quad \wedge \\
&\quad \mathit{generated}(A, I) = G_a \quad \mathsf{holdsAt} \quad T \quad \wedge \\
&\quad \mathit{demanded}(A, I) = D_a \quad \mathsf{holdsAt} \quad T \quad \wedge \\
&\quad G_a > D_a \qquad \wedge \qquad P_a = G_a - D_a
\end{aligned}
$$

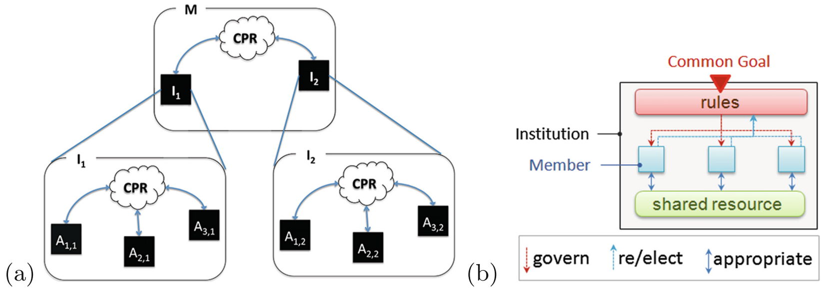

Fig. 1. (a) Nested CPR institutions and opaqueness of holons; (b) Conceptual model of a generic institution

The first EC axiom states that, in institution I, an agent A occupying the *prosumer* role is obliged to provision everything it generates to the common pool (the provision rule in I is *all*). In the second axiom, the provision rule is *excess*, so the obligation is only to provision the excess difference between what A generated and what it needed.

Note that in an open, decentralised system with autonomous components, as far as the institution is concerned, the prosumers are black boxes, and their 'internals' are unknown; therefore there are other rules to deal with incentives for compliance, monitoring and non-compliance, etc.

However, even within an institution, an *economy of scarcity* may occur when insufficient resources are generated to satisfy all prosumer demands. In this case, it would be beneficial to form 'alliances' with other institutions, and in times of excess it would contribute surpluses to a higher-order common pool, in the expectation of being allocated resources from that common pool in case of a shortfall later on. Note again that participation in the higher-level institution is subject to the mutual agreement of rules of engagement between the institutions; and that just as the prosumers were black boxes to their institutions, the institutions are essentially black boxes to the nested enterprise – the higher-level institution has no knowledge, or any need for any knowledge, of how the components choose to (self)-organise their own affairs (Fig. 1-a).

3 Institutions and Holons

It is the nesting of the rules of engagement at different levels of abstraction, and the opacity of components at each level, that suggests a relationship between multiple institutions as nested enterprises and holonic systems architectures. In this section we consider the convergence of institutions and holonics, which will yield the concept of holonic institutions.

3.1 Institutions: An Informal Overview

From a systemic perspective, an institution has a *well-defined goal* or objective, which it pursues by enforcing a *set of rules* on its *members*, or participants

(Fig. 1-b). For instance, in the context of electrical power sharing each community member may be entitled to receive a quota of available power for consumption at any one time; or different members may receive different quotas depending on the urgency of their consumption (e.g. medical facilities versus entertainment). Hence, institutions provide the necessary regulations and infrastructure for coordinating the actions of their members, which may otherwise diverge because of their *inherent dissimilarities* (e.g. in individual purposes and/or behaviours). In the absence of effective coordination, groups of non-identical members would most likely fail to achieve a common goal or compromise that would benefit all. An institution's purpose, rules, members and operational context may change over time requiring adequate *adaptations*.

An important question here is related to the manner in which the different functions and membership *roles* of an institution will be implemented. For instance, roles requiring more extensive insights or judgements may be assigned to human operators and performed over longer periods (e.g. redefine common goals and rules, based on knowledge and feedback), while more routine roles may be assigned to automated agents with reactive capacities (e.g. membership control, monitoring, policing and basic conflict resolutions).

3.2 Self-organising Electronic Institutions

The framework of dynamic norm-governed systems [1] defined three components: a specification of a norm-governed system; a number of changeable parameters, each with a range of values; and a stack of protocols detailing how to change the specification from one instance to another (i.e. change one parameter value for another). This effectively defined a kind of metric space with 'distances' between one specification instance and another. One way to define the protocol stack was to use an Action Language, such as the Event Calculus (as above). This also enabled constraints to be placed on the transition from one specification instance to another, for example on 'distance', but also some specification instances could be identified as non-normative and moving to them declared invalid.

This framework is very general: therefore in the class of dynamic norm-governed systems, we are interested in the sub-class in which the protocols formalise, in the Event Calculus, six of Ostrom's eight institutional design principles. This sub-class is referred to as self-organising electronic institutions. However, that work stopped short of formalising the eighth principle, and suggested further investigation of nested enterprises in several directions, including "the embedding of institutions within larger institutions, rather than the single layer model implemented here, to form the nested enterprises identified by Ostrom. ... [and the involvement of] third parties and other dependencies which can lead to other, more complex, supply chains" [20, p. 34]. We argue that this further investigation can be facilitated by using the principles of holonic systems.

3.3 Holonic Systems

In short, a *holonic system* (or *holarchy*) is composed of interrelated subsystems, each of which are in turn composed of sub-subsystems and so on, recursively, until reaching a lowest level of 'elementary' subsystems. As emphasised by Koestler [9], each such intermediary sub-system must play a dual role and be both: an autonomous whole controlling its parts; and a dependent part of a supra-system. This helps construct large systems with macro-goals from intermediary components able to achieve partial goals. There are several advantages that holonic structures provide for building viable complex systems [26]. These can be leveraged in applying holonic system principles to electronic institutions (and/or socio-technical systems), offering complexity management support by helping:

– institutions **scale** with the *number* and the *heterogeneity* of their members, since lower memberships in each holon put less strain on the institutional apparatus and decrease the level of internal diversity;
– to integrate institutions with **diverse goals**, since each of them only needs to be aware of the others' observable goals, state and negotiations, rather than of their internal details (e.g. rules and infrastructure);
– to improve an institution's local adaptation **reactivity**, while not directly impacting overall system **stability**, by the way in which the holonic structure modulates overall system dynamics and change propagation;
– system designers to **understand, analyse, simulate, adapt and predict** complex institutions, by allowing them to focus on a single holonic level at a time, with a reduced number of interrelated institutions.

4 Holonic Institutions

To benefit fully from these advantages, several important questions concerning holonic institutions have to be addressed:

– Q1: how to compose complementary or conflicting institutions?
– Q2: how to compose institutions at different scales, where each one can play the dual role of an autonomous institution and a semi-autonomous member?
– Q3: how to make holonic institutions adaptable, so that their goals can be achieved when changes occur in their environments, members, feedback on rule inefficiency, constraints from supra-institutions, or goal evolution?
– Q4: how to merge all the concerns above for constructing complex holonic institutions that can achieve their goals?

Figure 2-a depicts a generic conceptual model (abstract architecture) of holonic institutions to help address the questions above. In short, each holonic institution features two complimentary regulatory components implementing their dual roles. *Inward regulation* includes the internal rules, governance and adaptation functions for achieving a goal – as in Fig. 1-b; the difference being that this goal may diverge somewhat from the members' common goal since they agreed to join a supra-institution. *Outward regulation* merges, via conflict resolution and negotiation, the institution's own common goal with the (supra-)institutions' common goals. This results in the compromise goal that the

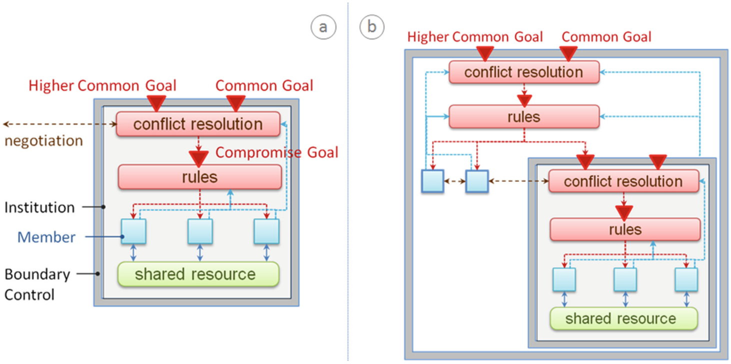

Fig. 2. (a) institution *holon* with dual role: inward/selfish &outward/transcendental; (b) supra-institution with several institutions/members

institution agrees to pursue. Each holonic institution is encapsulated within a *membrane* providing *membership-control* functions. At a high level of abstraction, this approach helps address institution composition questions (Q1 and Q2) (Figure 2-b). Institution adaptation relies on feedback from members and from the institution's evaluation of its goal achievement; it is propagated progressively from lower to upper holonic levels (Q3). This component-oriented design helps formalise, understand and analyse composite institutions, providing a key base for addressing the challenge of institutional complexity (Q4).

The above considerations provide a generic architectural overview on the manner in which *holonic institutions* can be constructed and maintained to address the aforementioned questions and achieve the advantages enabled by holonic principles. However, an important consideration is how Ostrom's seven other institutional design principles are impacted, in order to enable the eighth one – i.e. nested enterprises realised by holonic institutions. We will consider each design principle **(Px)** in turn.

(P1) Boundaries: *"who is and is not a member of the institutions should be clearly defined, as are the resources that are the subject of allocation"*. We propose to encapsulate each institution holon within a special-purpose container, or membrane, which helps isolate a holon's interior from the rest and separate between its internal resources and external environment. This reduces the holon's internal complexity as it only involves a 'manageable' sub-set of the entire system components; and, a predefined set of exchanges with its environment, controlled by its membrane (e.g. message filtering and aggregation). The membrane also exposes the holon's interfaces, allowing holonic institutions and members to appear identically to external observers. They include: the holon's goal; the feedback on purpose achievement; and the inter-holon negotiations, see Fig. 2.

(P2) Congruence: *"the rules should be congruent with the prevailing local environments (including the profile of the members themselves)"*. This rule will have a decisive impact on the overall shape of the holarchy – i.e. how members group into institutions, and institutions into supra-institutions, recursively. This will impact the size of each holon, as an institution's 'manageable' size will depend on the highest degree of divergence that can be supported over its group members. As stated previously, institutions are about coordinating divergent populations in order to achieve globally advantageous compromises. Hence, each member has to diverge somewhat from its selfish purposes and behaviours in order to benefit from the institution (cf. P6 minimal rights). An institution's internal divergence would have to be limited so as to allow for compromises that are both: sufficiently specific to be effective for achieving the group's purpose; and, sufficiently general to be acceptable to group members. Therefore, successful institution holons are more likely to be obtained by grouping members and institutions with most similarities (e.g. rather than via geometric borders).

(P3) Participation: *"those individuals who are affected by the collective choice arrangements should participate in their selection"*. To apply this rule, each institution holon must be able to define a common goal as an aggregate of the goals of its members, and use this aggregate goal when participating in the selection of higher-level rules for its supra-holons. Electing representative members to carry-out such negotiations may also be considered. Priorities must be set when participating in several supra-institutions with conflicting goals.

(P4) Monitoring: *"compliance with the rules should be monitored by the members themselves, or by agencies appointed by them"*. To apply this rule, each institution holon must be able to provide an aggregate estimate of the degree to which it has achieved its goal, or complied to its supra-institutions rules, based on estimates from its internal institution holons, or members.

(P5) Sanctions: *"graduated sanctions should ensure that punishment for non-compliance is proportional to the seriousness of the transgression"*. To apply this rule, each institution holon must be able to translate, proportionally, external sanctions for the institution to specific sanctions for its members.

(P6) Conflicts: *"the institution should provide fast, efficient and effective recourse to conflict resolution and conflict prevention mechanisms"*. Each institution holon must be able to detect and resolve conflicts between the common goal of its supra-institution (external) and the own common goal of its individual members (internal). Figure 2 depicts this via a specific *conflict resolution* component, which computes compromises between external and internal goals. These compromises are first *negotiated* with the other members of the supra-institution and then forwarded to the holon's internal members.

(P7) Minimal recognition of rights to self-organise: *"the rights of appropriators to form their own institutions are not challenged by external authorities"*. Holonic institutions will generally have fewer degrees of freedom in order to be integrated within a higher-level supra-institution, i.e. its autonomy may be limited because more specification instances become non-normative or invalid. This was actually already the case for prosumer members of base level institutions (cf. congruence). The acceptability of constraints and restrictions on rule formation would depend on the benefits expected from joining the institution. The more a member's own goal diverges from an institution's common goal, the bigger the required compromise and so the amount of autonomy that a member will have to surrender for staying in that institution. Similarly, an institutional holon needs to give up an amount of autonomy that is proportional to the difference between its supra-institution's goal (global compromise) and its internal goal. The exact proportionality can also be modulated by various configurations of the sanctions (P5) and boundary control (P1) rules.

In fact, the move towards holonic institutions provides for a much finer-grained separation of rights and powers. It is *not* that specific instances of rules themselves cannot be challenged: it is the *right* to form the institution, and to self-organise its rules, that is at issue. This is the fundamental issue in design principle P7, and is strikingly exposed by holonic systems thinking.

5 Case Study for Community Energy Systems

In this section we apply the concept of holonic institutions to the Smart Grids case study in subsect. 2.3. First we discuss Smart Houses as the basic holonic unit. We then introduce the idea of decentralised community energy systems (dCES) as the basic holonic institution unit. This leads to an analysis of multiple dCES as nested enterprises forming holonic institutions, and of the various Smart Grids agencies leading to polycentric self-governance and adaptivity.

5.1 Smart Houses

Smart Houses include technical systems that aim to automate *residential services* – e.g. safety and security, home entertainment, control of heating, ventilation and air conditioning – in order to improve owner comfort and experience. Since these systems operate in a social context, Smart Houses become socio-technical systems where several objectives, both technical and social, must be met.

Smart Houses do not operate in isolation and must integrate 'smoothly' within larger socio-technical systems – e.g. smart cities and electric grids. Several authorities with diverse interests and objectives operate at these levels, including city representatives and power grid operators. While each Smart House must remain largely autonomous and pursue its owners' objectives, it must also yield some of its autonomy in order to comply with the more global socio-technical systems that it joins for achieving a broader common purpose.

5.2 Decentralised Community Energy Systems

A Community Energy System (CES) is an energy generation, distribution and storage system involving local community ownership and participation. Generally, the differentiation between a nationally and community operated system is the boundary of autonomy – where responsibility for network specification and operation switches from the grid operator to the CES operator [25].

A *decentralised* CES (dCES), illustrated in Fig. 3, is a network of geographically co-located Smart Houses installed with small-scale renewable sources like

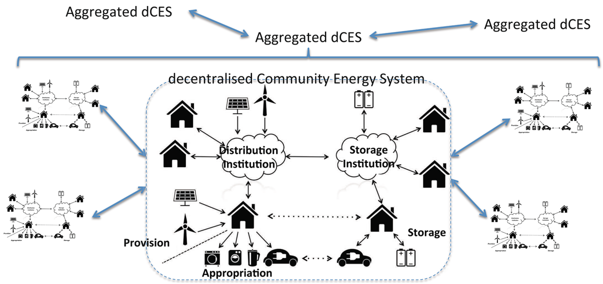

Fig. 3. Decentralised Community Energy System (dCES)

photovoltaic (PV) cells or micro wind turbines. At this base level, we assume there is no enterprise-owned Combined Heat and Power (CHP) plant or other large-scale generation: everything is generated in-house (literally), controlled and operated by the residents of Smart Houses. Storage can be provided by in-house batteries or, looking farther ahead, electric vehicles. A group of dCES can be aggregated into a larger institution, as previously discussed.

5.3 Holonic Institutions

If we think of an individual Smart House as a single holon, then we can create institutions at the base level by forming a dCES comprising multiple Smart Houses with a set of institutional rules meeting Ostrom's design principles. This allows for a wide range of institutional types. For example, we could have one type of institution whose energy distribution is based on the formalisation of social relationships, such as legitimate claims [19,22], and another type which is primarily market-oriented. Assuming that the minimal membership requirement is met – i.e. to have installed some renewable energy generation and/or micro-storage facility – then two (of many) types of dCES are summarised in Table 1.

In [25], four types of CES were identified: multi-home energy schemes, as suggested here; local energy schemes; district schemes with enterprise collaboration; and, district scheme with large generation. The different types were distinguished according to their ownership model, generation and storage facilities, and grid relationship. This latter could be *grid forming*, if the system operates pre-dominantly independently of the (national) grid (for example in terms of frequency and voltage control); *grid following*, which maintain voltage and frequency using the grid as reference; and, *grid supporting*, if they operate in parallel with the grid for the purposes of importing and exporting power.

It might be that the objective of a type-1 dCES would be grid-forming, while the objective of of a type-2 dCES would be grid-supporting, with the intention to export power through (in the UK) a FIT (feed-in tariff) scheme. Therefore, we can see how a dCES could form different institutional relationships with larger generation schemes, such as CHP plants, to form institutions at a larger scale. The larger scale institutions have different objectives, some technical (e.g.

Table 1. Institutional rules for two types of dCES

dCES Type 1: Social relationship-oriented	
Ostrom principle	Implementation
P1. Membership	By invitation
P2. CPR rules	Provision and appropriation according to legitimate claims [19]
P3. Participation in rule selection	One member one vote
P4. Monitoring	SmartMeter[a]
P5. Sanctions	Diminished claims
P6. Conflicts	Alternative dispute resolution
dCES Type 2: Market-oriented	
Ostrom principle	Implementation
P1. Membership	By subscription
P2. CPR rules	Market-based (e.g. auction)
P3. Participation in rule selection	Enterprise appointed management board
P4. Monitoring	SmartMeter
P5. Sanctions	Cash fines
P6. Conflicts	Court hearing

[a] With caveats, as discussed in Sect. 6

voltage and frequency control), some economic (e.g. import and export of power), and some political (e.g. meeting low-carbon targets). Critically, we can see these larger-scale institutions being realised in the framework of Fig. 2-b.

5.4 Polycentric Self-Governance

From a wider perspective, a Smart Grid is, like the water basins of California studied by Ostrom, composed of numerous actors and agencies with different ownership models – e.g. private individual, mutual cooperative, private enterprise, national infrastructure and regulator. Table 2 identifies a number of institutions involved in dCES, together with their associated *common goals* as indicated in Fig. 2.

It is well-known that managing critical infrastructure, like a national energy generation, transmission and distribution network, will necessarily involve multiple agencies with differing (possibly competing or even conflicting) interests, effectively creating a kind of overlay network of relational dynamics which also needs to be resolved. Furthermore, there is some, not always well-understood, inter-connection of public and private ownership that makes the overall system both stable and sustainable.

Therefore, in analysing any such complex system, it is critical to identify the agencies and determine their institutional *common goals* – what each agency (through its institution) is trying to achieve or maintain, by coordination with

Table 2. Actors/Agencies in dCES

Agency (Institution)	Common purpose/typical functions
Administration	CPR management – operate the servers running CPR Apps – compute the resource allocation – apply membership rules
Appropriators	Meet production/storage power goals – provision and appropriate energy (generation and storage) – investment strategy
Service providers	Infrastructure and equipment – grid connectivity, voltage and frequency control – installation and maintenance of micro scale generation and storage facilities – market access (e.g. FIT)
Ombudsman/courts	Dispute resolution – legal representation – negotiation, mediation and arbitration
Regulators	Consumer protection – protect present and future consumer's interests – meeting national and international policy goals
Citizens' advocacy	Accountability, pressure, special interests – represent environmental/green energy interests
Policy officials	Regulations (at multiple scales) – policy drafting – advice and calibration of CPR/CES rules
App entrepreneurs	Software service development – SmartMeter Apps

other institutions and by the decision-making of its members. Such analysis makes it possible to understand the 'ecosystem' of institutions and how they fit together as collaborators or competitors, based on the nature of their goals and the scope of their influence.

In this way, we believe that all of the institutions (nested enterprises) identified in Table 2 can be organised in a holonic manner. The outcome is twofold. Firstly, that it supports polycentric self-governance at all scales of the system, and in particular supports *subsidiarity* (the idea that problems are solved as close to the local source as possible). Secondly, it encourages the institutions to recognise their role in the overall 'scheme of things' in relation to institutions at the same, higher and lower levels. This is a key requirement for *adaptive* institutions [21] and this establishment of *systems thinking* as a commonplace practice within any one institution is what we may refer to as *institutionalised holonics.*

6 Related and Future Work

There are many theories and tools for organisation in multi-agent systems, including MOISE (an organisational model for multi-agent systems) [8], OMACS (organisational model for adaptive computational systems) [3], LAO (logic and organisations) [4], LGI (law governed interaction) [12], electronic institutions [7], and others, but none of these works takes holonic design into account. The issue of multiple interacting institutions has been addressed in [2,18], but only peripherally (at best) consider the concepts addressed here: norm-governed institutions, Ostrom's institutional design principles, holonic system architectures, and polycentric self-governance. Holonic multi-agent systems (HMAS) have been proposed in [23] yet *not* applied to social systems or electronic institutions.

Equally, several works use holonic design patterns to develop technological artefacts and complex systems [29], including traffic control [5], manufacturing plants [28] and (of course) Smart Grids [6]. However, we are not aware of any work that explicitly represents institutional or organisation concepts inside the holon, and reasons with these, with respect to its common goal.

There is much valuable work on a system of systems approach to complexity and self-organisation [27]. To the best of our knowledge, though, the present paper is the first work that has attempted to converge the hitherto disjoint works on self-organising institutions (based on conventional rules formalising Ostrom's design principles) and holonic architectures: i.e. to address both the functional and structural properties of complex CPRs in the context of a single unified framework and its application to a complex system like a Smart Grid.

Evidently, the proposal of holonic institutions and the case study presented in this paper are conceptual rather than actual. In further work, we plan to formalise and implement the concepts both in multi-agent simulation and a Smart Grid testbed, in particular to understand the relationship between structure and macro-level properties such as robustness, stability, resilience and sustainability.

However, in modelling and simulating socio-technical systems of this kind, there are other dimensions to consider. One is the relationship between people and institutions and the incorporation of processes from dynamical social psychology (e.g. [13]) into this framework. Another is the effect that some political/regulatory decisions may yet have on the evolution of the Smart Grid. If the so-called SmartMeter is unbundled (separating the platform from the the grid itself), as advocated by [24], this will have a telling impact on the Smart Grid 'institutional ecosystem'. Modelling this process is essential for understanding and responding to a new wave of innovation (driven by the App Entrepreneur agencies in Table 2) in a constructive and meaningful way.

7 Summary and Conclusions

This paper is situated within a broader research programme concerned with the formalisation and operationalisation of Ostrom's institutional design principles to engineer self-* properties for management and control of complex open systems. Specifically, it has focused on the formalisation of the eighth principle: *"For CPRs that are part of larger systems, nested enterprises"*. Since this principle relates more to structure rather than function, it has proved difficult to formalise in a declarative specification, like Principles 1–6, for electronic institutions [20]. Accordingly, the approach that has been proposed in this paper has been based on *structures* and *architectures* rather than rules.

The contribution of this paper is therefore threefold. By converging previously disjoint approaches to the design of complex open systems, one based on electronic institutions [19,20] and the other based on holonic architectures [6], the paper has contributed:

- a critical analysis of Ostrom's eighth institutional design principle for electronic institutions and socio-technical systems;
- an innovative proposal for *holonic institutions*, whereby institutions can be composed and de-composed as nested enterprises, enabling multi-scale polycentric decision-making to be established in the ecosystem of organisations;
- a case study in using holonic institutions for polycentric self-governance in community energy systems (smart grids).

This is, of course, only a first step in developing, demonstrating and applying such concepts. However, if successful, the ultimate contribution of this research could be to enhance polycentric theory, as a branch of political science, with the technology and tools to both analyse and design complex, multi-scale socio-economic, socio-political and socio-technical systems. These in turn would help address complex, multi-scale ecological challenges, such as climate change, just as Ostrom proposed [15].

Acknowledgements. Jeremy Pitt was partially supported by the UK EPSRC Grand Challenge project *The Autonomic Power System* (EP/I031650/1).

References

1. Artikis, A.: Dynamic specification of open agent systems. J. Logic Comput. **22**(6), 1301–1334 (2012)
2. Cliffe, O., De Vos, M., Padget, J.: Specifying and reasoning about multiple institutions. In: Noriega, P., Vázquez-Salceda, J., Boella, G., Boissier, O., Dignum, V., Fornara, N., Matson, E. (eds.) COIN 2006. LNCS (LNAI), vol. 4386, pp. 67–85. Springer, Heidelberg (2007)
3. DeLoach, S.: OMACS: a framework for adaptive, complex systems. In: Dignum, V. (ed.) Multi-Agent Systems Semantics and Dynamics of Organisational Models. IGI Global, Hershey (2009)
4. Dignum, V., Dignum, F.: A logic of agent organizations. Logic J. IGPL **20**(1), 216–283 (2012)
5. Fischer, K.: Holonic multiagent systems – theory and applications. In: Barahona, P., Alferes, J.J. (eds.) EPIA 1999. LNCS (LNAI), vol. 1695, pp. 34–48. Springer, Heidelberg (1999)
6. Frey, S., Diaconescu, A., Menga, D., Demeure, I.: A holonic control architecture for a heterogeneous multi-objective smart micro-grid. In: IEEE 7th International Confrence on Self-Adaptive and Self-Organizing Systems (SASO), pp. 21–30 (2013)

7. García-Camino, A., Noriega, P., Rodríguez-Aguilar, J.: Implementing norms in electronic institutions. In: Proceedings of the Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), pp. 667–673. ACM Press, NewYork (2005)
8. Hübner, J., Sichman, J., Boissier, O.: Developing organised multiagent systems using the MOISE$^+$ model. IJAOSE **1**(3/4), 370–395 (2007)
9. Koestler, A.: The Ghost in the Machine. Hutchinson Publisher, London (1967)
10. Kohler, T., Steghöfer, J.P., Busquets, D., Pitt, J.: The value of fairness: trade-offs in repeated dynamic resource allocation. In: Proceedings 8th IEEE Conference on Self-Adapting and Self-Organising Systems (SASO) (2014)
11. Kowalski, R., Sergot, M.: A logic-based calculus of events. New Gener. Comput. **4**, 67–95 (1986)
12. Minsky, N.: Law-Governed Interaction (LGI): A Distributed Coordination and Control Mechanism. Rutgers University (2005)
13. Nowak, A., Vallacher, R., Kus, M., Urbaniak, J.: The dynamics of societal transition: modeling non-linear change in the polish economic system. Int. J. Sociol. **35**, 65–68 (2005)
14. Ostrom, E.: Beyond markets and states: polycentric governance of complex economic systems. In: Grandin, K. (ed.) Les Prix Nobel, The Nobel Prizes 2009, pp. 408–444. Nobel Foundation, Stockholm (2010)
15. Ostrom, E.: Thinking about climate change as a commons. In: 15th Annual Philip Gamble Memorial Lecture, pp. 1–34. UMass Amherst (2011)
16. Ostrom, E., Ahn, T.: Foundations of Social Capital. Edward Elgar Pub, Cheltenham (2003)
17. Ostrom, E.: Governing the commons. Cambridge University Pressm, Cambridge (1990)
18. Patel, J., Teacy, W., Jennings, N., Luck, M., Chalmers, S., Oren, N., Norman, T., Preece, A., Gray, P., Shercliff, G., Stockreisser, P., Shao, J., Gray, W., Fiddian, N., Thompson, S.: Agent-based virtual organisations for the grid. Multiagent Grid Syst. **1**(4), 237–249 (2005)
19. Pitt, J., Busquets, D., Macbeth, S.: Distributive justice for self-organised common-pool resource management. ACM Trans. Auton. Adapt. Syst. **9**(3), 14 (2014)
20. Pitt, J., Schaumeier, J., Artikis, A.: Axiomatisation of socio-economic principles for self-organising institutions: concepts, experiments and challenges. ACM Trans. Auton. Adapt. Syst. **7**(4), 39:1–39:39 (2012)
21. RCEP: 28th report: Adapting institutions to climate change. Royal Commission on Environmental Protection, The Stationery Office Limited (2010)
22. Rescher, N.: Distributive Justice. Bobbs-Merrill, New York (1966)
23. Rodriguez, S., Gaud, N., Hilaire, V., Galland, S., Koukam, A.: An analysis and design concept for self-organization in holonic multi-agent systems. In: Brueckner, S.A., Hassas, S., Jelasity, M., Yamins, D. (eds.) ESOA 2006. LNCS (LNAI), vol. 4335, pp. 15–27. Springer, Heidelberg (2007)
24. Sanduleac, M., Eremia, M., Toma, L., Borza, P.: Integrating the electrical vehicles in the smart grid through unbundled smart metering and multi-objective virtual power plants. In: 2nd IEEE PES International Conference Innovative Smart Grid Technologies, ISGT Europe 2011, pp. 1–8 (2011)
25. SESIG: Smart Energy Special Interest Group: The role of community energy systems in the UK resilient energy supply. Technical report, Innovate UK (Technology Strategy Board) (2013)
26. Simon, H.: The architecture of complexity. Proc. Am. Philos. Soc. **106**(6), 467–482 (1962)

27. Steghöfer, J.P., Anders, G., Siefert, F., Reif, W.: A system of systems approach to the evolutionary transformation of power management systems. In: GI Jahrestagung, pp. 1500–1515 (2013)
28. Ulieru, M., Brennan, R., Walker, S.: The holonic enterprise: a model for internet-enabled global manufacturing supply chain and workflow management, integrated Manufacturing Systems. Integr. Manuf. Syst. **13**(8), 538–550 (2002)
29. Valckenaers, P., Brussel, H.V., Holvoet, T.: Fundamentals of holonic systems and their implications for self-adaptive and self-organizing systems. In: IEEE SASO Workshops (SASOW) (2008)

http://www.springer.com/978-3-319-25419-7

Coordination, Organizations, Institutions, and Norms in Agent Systems X
COIN 2014 International Workshops, COIN@AAMAS, Paris, France, May 6, 2014, COIN@PRICAI, Gold Coast, QLD, Australia, December 4, 2014, Revised Selected Papers
Ghose, A.; Oren, N.; Telang, P.; Thangarajah, J. (Eds.)
2015, X, 269 p. 66 illus. in color., Softcover
ISBN: 978-3-319-25419-7