

THE ALGORITHMIC GOVERNANCE OF COMMON-POOL RESOURCES

By Jeremy Pitt and Ada Diaconescu

As we saturate our everyday environment with computing and communication technologies, we can increase our capacities for successful collective action if we devise system infrastructures to support self-organization, self-management and pro-social behavior. Elinor Ostrom's institutional design principles for managing common-pool resources provide a valuable template for designing effective Internet-based applications for algorithmic self-governance.

Introduction: Resource Allocation in Open Systems

Using a methodology called sociologically inspired computing,[1] researchers are now attempting to solve engineering problems by developing “formal models of social processes.” This entails examining how people behave in similar situations and, informed by a theory of that behavior grounded in the social sciences, developing a formal characterization of the social behavior (based on the theory) using mathematical and/or computational logic. This logical specification then provides the basis for the specification of an algorithmic framework for solving the original problem.

In networks that function as open systems, for example, a significant challenge is how to allocate scarce resources.

This is a vexing challenge because open computing systems and networks are formed on the fly, by mutual agreement, and therefore they may encounter situations at run-time that were not anticipated at design-time. Specific examples include ad hoc networks, sensor networks, opportunistic and vehicular networks, and cloud and grid computing. All these applications have at least one feature in common: the system components (henceforth referred to as *agents*) must somehow devise a means to collectivize their computing resources (processor time, battery power, memory, etc.) in a common pool, which they can then draw upon in order to achieve their individual goals in a group (or as a group) that they would be unable to do if they each functioned in isolation.

However, open systems face serious challenges in coordinating agents because there is no centralized controller-agent that is compelling other agents in the system to behave in a certain way with regards to the provision and appropriation of resources. Furthermore, all agents may be competing for a larger share of the common pool, and may therefore not comply with the requirements for “correct” (pro-social) behavior. For example, they may appropriate resources that they were not allocated, or they may appropriate resources correctly but fail to contribute expected resources (a phenomenon known as “free riding”).

The Tragedy of the Commons

So, applying the first step in the methodology of sociologically inspired computing, the question is: How do people collectively allocate shared common-pool resources when there is no centralized controller “dictating” what the allocation should be? This prompts two further questions: 1) How do they solve problems such as free riding and other examples of anti-social behavior? 2) How do people allocate common-pool resources in a way that is considered *fair* by the individuals involved (in some sense of the word “fair”), and is *sustainable in the long-term* (i.e., the self-renewing resource, like a forest or fishery, is properly managed and maintained so that it is not over-used)?

One analysis of this problem, called the *tragedy of the commons*, suggests that there is *no* internal solution to this problem. According to biologist Garrett Hardin, people will inevitably deplete (exhaust) common-pool resources in the short term even if that is in no one’s interest in the long-term. Many people assume that the only way to ensure that such resources are maintained is through externalized oversight by some centralized body (e.g., government) or through privatization. These solutions are not, of course, available to engineers of open computing systems.

Ostrom’s Self-Governing Institutions

Although some economists believe that people will inevitably deplete common-pool resources to which they have access, the empirical data of hundreds of case studies suggests that other outcomes are possible. For example, based on extensive fieldwork, from water irrigation systems in Spain to fisheries in Switzerland and forests in Japan, economist and political scientist Elinor Ostrom observed that actually people tend to cooperate in such collective action situations, not only to avoid depleting the resource, but to actively maintain it, even over the course of generations.[2] Ostrom was awarded the Nobel Prize for Economic Science in 2009 for her extensive fieldwork and theoretical innovation in demonstrating the feasibility of managing common-pool resources.

The essence of the many counter-examples is this: it turns out that people are very good at *making stuff up*. In particular, people are very good at making up rules to deal with novel circumstances and challenges. Without the ability to make up rules, there would be no playing of games, for example; nor would people be able to improvise coordinated responses to emergencies. Similarly, Ostrom observed in many collective action situations, people make up rules to *self-determine* a fair and sustainable resource allocation. People voluntarily agree to abide by and regulate their behavior. Notably, these are not immutable physical laws; they are social conventions that people can and sometimes do break – either by accident, necessity or (sadly) sheer malice.

The invention of conventional rules (and their rationalization and stabilization into what Ostrom called *institutions*) is a necessary condition for preserving resources over time, but it is not a sufficient condition. On some occasions when communities develop an institution to manage their affairs, the resource is successfully sustained, but sometimes it is not. Addressing the requirement to *supply* self-governing institutions for enduring common-pool resource management, Ostrom proposed eight institutional design principles:

1. *Boundaries*: who is and is not a member of the institutions should be clearly defined – along with the resources that are being allocated;
2. *Congruence*: the rules should be congruent with the prevailing local environments (including the profile of the members themselves);

3. *Participation*: those individuals who are affected by the collective choice arrangements should participate in formulating and adopting them;
4. *Monitoring*: compliance with the rules should be monitored by the members themselves, or by agencies appointed by them;
5. *Proportionality*: graduated sanctions should ensure that punishment for non-compliance is proportional to the seriousness of the transgression;
6. *Conflicts*: the institution should provide fast, efficient and effective recourse to conflict resolution and conflict prevention mechanisms;
7. *Autonomy*: whatever rules the members agree to govern their affairs, no external authority can overrule them;
8. *System of systems*: multiple layers of provisioning and governance should be nested within larger systems.

A meta-review has confirmed these principles with only minor adjustments.[3]

A Formal Characterization as Electronic Institutions

Elinor Ostrom's research provides a theory of how people can solve the collective-action problem of common-pool resource allocation. To use this theory as a basis for engineering solutions in open computing systems, three related questions must be addressed: 1) Can the theory of self-governing institutions be given a formal characterization in computational logic? 2) Can the computational logic specification be given an algorithmic interpretation that can be used to implement a self-organizing *electronic* institution? 3) Can the agents in a self-organizing electronic institution be designed according to Ostrom's eight principles so as to successfully manage and sustain a common-pool resource?

Pitt, Schaumeier and Artikis give a positive answer to all three questions.[4] The first six of Ostrom's principles were each axiomatized in first-order logic using the Event Calculus, a language used in Artificial Intelligence to represent and reason about actions and the effects of actions. This axiomatic specification was then converted into Prolog and queried as a logic program: i.e. *the specification is its own implementation*. As such, the set of clauses comprising the logic program constitutes an algorithmic specification for self-governance. Finally, the implementation was tested in a multi-agent resource allocation system that allowed clauses for each principle to be individually included in successively more complex experiments. The results showed that the more principles that were included, the more the agents (as members of the institution) were able to sustain the resource and maintain a high membership.

But Is It Fair? Distributive Justice and the Canon of Legitimate Claims

These experiments demonstrated that Ostrom's institutional design principles for managing enduring common-pool resources could provide the basis for achieving sustainable resource allocation in open computing systems. One complication, however, is that certain elements of human social systems cannot necessarily be represented in the logic of electronic "social" systems. For example, in establishing the congruence of the provision and appropriation rules to the prevailing state of the environment (Principle 2), a software designer might assume that if Principle 3, requiring user participation in

making rules, were in place, then those affected by the rules would select rules that were intrinsically or implicitly “fair.” This assumption cannot be made in electronic networks whose components are without any understanding of a concept of ‘fairness,’ however.

To address this issue, Pitt, Busquets and MacBeth [5] suggest applying the methodology from another theory from the social sciences, the theory of distributive justice articulated by the philosopher Nicholas Rescher.[6] Rescher observed that distributive justice had been held, by various sources, to consist of treating people wholly or primarily according to one of seven canons (established principles expressed in English). These canons consist of equality, need, ability, effort, productivity, social utility and supply-and-demand. However, these canons each have different properties and qualities, and they therefore speak to many different (and possibly inconsistent) notions of utility, fairness, equity, proportionality, envy-free conviviality, efficiency, timeliness, etc.

Rescher’s analysis showed that each canon, taken in isolation, was inadequate as the sole criterion of distributive justice. He proposed instead that distributive justice could be represented by the *canon of claims*, which consists of treating people according to their legitimate claims, both positive and negative. Then the issue of “Which is the preferred canon of distributive justice?” can be displaced by questions such as: “What are the legitimate claims in a specific context, for fairness? How can plurality be accommodated? How can conflicts be reconciled?”

Pitt, Busquets and MacBeth implemented another multi-agent system testbed and conducted another set of experiments to explore resource allocation in an economy of scarcity.[5] (This scenario is defined as one in which there are insufficient resources at any time-point for everyone to have what they demand, but there are sufficient resources over a succession of time-points for everyone to get enough to be ‘satisfied’). In this testbed, each of the canons (if it was relevant in this context) was represented as a function that computed an ordering of the agents requesting resources. To address the plurality of claims, the functions were then used in a weighted Borda Count – a voting protocol that computes an overall rank order and is more likely to produce a consensus outcome rather than a simple majoritarian outcome. To reconcile conflicts among claims, the agents themselves decided the weight to be associated with each canon in prioritizing the agents’ claims.

The results showed that a group of agents, in an electronic institution based on Ostrom’s principles, could self-organize a distribution of resources using the canon of legitimate claims such that it was *fair over time*. That is, while at any one time-point the resource allocation might be very unfair (using a well-known and often-used fairness metric, the Gini index), a group could nonetheless achieve allocations that were very fair over a series of time-points. The distribution could also be made fairer than alternative allocation schemes based on random assignment, rationing or strict queuing.

Socio-Technical Systems

The formalization and implementation of social processes, such as Ostrom’s institutional design principles and Rescher’s theory of distributive justice, provide an algorithmic basis for governance of common-pool resources in electronic social systems. These are not models of how human social systems work – but nor are they intended to be. Instead of asking if these are testable models with predictive or explanatory capacity (adequacy criteria for this are included in the methodology [1]), a more pertinent follow-

up question is: Can this formal approach to algorithmic self-governance be injected into open *socio-technical* systems – i.e., systems in which human participants interact with an electronically saturated infrastructure, or with each other through an electronically-mediated interface, in trying to exploit, and sustain, a common pool resource?

Here are three examples in which algorithmic self-governance could be usefully applied in socio-technical systems: decentralized community energy systems, consensus formation in open plan offices, and ‘fair’ information practices in participatory sensing applications.

1. In a *decentralized community energy system*, a group of geographically co-located residences may be both producers and consumers of energy. For example, the residence may have installed photovoltaic cells, small wind turbines or other renewable energy source; and the residence occupants have the usual requirements to operate their appliances. Instead of each residence generating and using its own energy, and each suffering the consequences of over- or under-production, the vicissitudes of variable supply and demand could be evened out by providing energy to a common-pool and computing a distribution of energy using algorithmic self-governance. Furthermore, excessive demand, which would otherwise lead to a power outage, could be pre-empted by synchronized collective action in reducing consumption.

2. Similarly, an *open plan office* is a working environment that requires people to share a common space. However, a violation of conventional rules determining what is (and is not) acceptable behavior can cause instances of incivility which, if untreated, can lead to problems of escalating retaliation, a demoralized or demotivated workforce, staff turnover, and other problems. We have developed a prototype system in which we regard the (intangible) “office ambience” as a pooled resource which the office occupants can deplete by anti-social behavior and re-provision by pro-social behavior. The system interface supports consensus formation by enabling the office-workers themselves to determine what is (and is not) anti-social behavior, and supports them in detecting violations, issuing apologies and encouraging forgiveness. This is an instantiation of Ostrom’s third principle – that those affected by collective choice arrangements should participate in their selection. Ostrom’s fifth and sixth principles – dealing with the system of conflict prevention and resolution – should encourage pro-social behavior.

3. *Participatory sensing applications* are innovative systems that aggregate and manipulate user-generated data to provide a service. A typical example is taking users’ mobile phone location and acceleration data to infer traffic density and so provide a transportation advice service. However, in many of these applications, the generators of the data are not the primary beneficiaries, and furthermore, there are severe privacy concerns over who has access to this data, how long it is stored, and what is used for. An alternative approach is to regard this user-generated data as a *knowledge commons*, and regulate access through self-determined rules, and so achieve a “fair” return of service for user-generated data.

Adaptive Institutions and Algorithmic Governance: The Way Forward

Studies in technology and law have often referred to the *law lag*, in which existing legal provisions are inadequate to deal with a social, cultural or commercial context created by rapid advances in information and communication technology (ICT).[7]

We can reasonably refer to a similar phenomenon of “institution lag”, whereby the rate of technological advance far outstrips the ability of traditional institutions to adapt fast enough to track the activity it was intended to regulate. Yet adaptive institutions have been identified as a critical tool in addressing environmental challenges such as climate change and sustainability.[8]

The challenge of devising effective algorithmic governance has a lot to do with scale. We can observe that, at the micro-level, human participants are able to self-organize and self-adapt by playing various roles, but at the macro-level, the emergent outcomes of unrestricted self-organization may be ineffective or undesirable (e.g., it may result in a tragedy of the commons).

We believe that more desirable macro-outcomes may be achieved by introducing a meso-level of governance: a rule-based, ICT-enabled algorithmic framework for self-governance that is designed to assure that whatever emerges at the macro-level represents the self-identified best interests of the community’s majority. The ultimate result would be to create more flexible institutions that could adapt more quickly to rapid societal changes. Since such rapid societal changes are being *caused* by ICT, it makes sense that the rapid adaptation required may be best *enabled* by ICT. Indeed, this may be the only feasible approach.

The ICT-enabled framework would provide an interaction medium that inherently implements Ostrom’s rules, enabling participants to self-organize into “fair” institutions (avoiding the *tragedy of the commons*) and to self-adapt such institutions to contextual changes (avoiding the *institution lag*). Such ICT framework should enable participants to perform critical activities, such as defining community rules for resource sharing, boundary definitions and non-compliance sanctions. It should also provide core automatic functions that facilitate the participant’s tasks, including for instance: managing membership based on boundary definitions; evaluating participant compliance with rules and applying sanctions; ensuring protection from external intrusion and interference; and provisioning comprehensible feedback on emerging results such as “fairness,” at both micro and macro levels, which is critical for efficient rule adaptation. Finally, such a system must ensure essential properties such as overall stability, robustness and resilience, while preserving crucial social concepts like privacy, safety and security.

In this context, the meso-layer ICT framework is vital in helping to deliver the desired outcomes. This is why a platform like Open Mustard Seed (see Chapter XXX), which offers designers at least the opportunity to strike the right balance between continuity and stability on the one hand, and adaptivity and responsiveness on the other, is crucial if algorithmic governance of common-pool resources, and other forms of collective action, are to be successful.

At this stage, of course, there is much that we do not know. For instance, the ICT system’s *scalability* is an important concern. Here, scale relates to the total number of participants; the level of heterogeneity in targeted environments and participant profiles; the number of societal interests considered and perhaps also their cultural backgrounds; and, the incidence of conflicts among intersecting heterogeneous groups. Achieving and maintaining macro-objectives in a large-scale system composed of autonomous self-adaptive agents and situated in a continuously changing environment, will require a trans-disciplinary investigation across the social and computational sciences.

A common feature observable in most (or all?) natural systems of similar scales and dynamics, such as living organisms, societies or economies, is their reliance on *holonic* organizations (see Chapter YYY). As first described by Arthur Koestler in the 1960s, each system element is both an autonomous entity pursuing its own objectives and controlling its internal resources as well as an element nested within a higher-level organization and contributing to its higher-level objectives by following its control commands. Recursively composing elements in this manner results in a holonic organization, or “holarchy” – a hierarchy in which each element is both autonomous yet contained within higher-level structures.

A holarchy seems essential for managing scalability issues because the structure enables problems to be detected and dealt with in isolation, at the lowest possible level, without disrupting the larger system. The holonic structure also ensures that both micro (individual) and macro (community) objectives are met concomitantly.

Successfully delivering such systems would directly satisfy Ostrom’s eighth principle, i.e., a self-governing system of systems. But one of the critical difficulties here is the implementation of each community’s “dual nature” as both an autonomous community with its own objectives and fairness rules and as a participant in a larger community with higher-level objectives and equity goals. This dualism reflects the built-in tensions of any individual, who naturally pursues personal objectives (selfish nature) while respecting larger community objectives (societal or transcendental nature).

Unresolved Issues

There are a number of issues that remain unresolved in devising systems of algorithmic self-governance, however. One involves the various conflicts that may occur when members belong to several communities with incompatible notions of fairness. Once these challenges are addressed theoretically, the ICT framework could in principle implement the necessary infrastructure and mechanisms for ensuring that the targeted system could self-organize into a holonic structure featuring the desired properties.

The “social ergonomics” of self-governance platforms is another important aspect that will need to be evaluated and refined. Notably, even if the macro-objectives emerging at any one time are fair with respect to a society’s common good, and even if fairness is ensured in the long-term for each individual, this will not necessarily imply an acceptable *experience* for each individual in that society. For instance, while *change* may be essential for ensuring fairness in a dynamic environment, change may also cause considerable distress and discomfort to individuals experiencing it. From an individual’s perspective, a relatively “unfair” state of affairs, in which they can comfortably survive in more or less stable circumstances, may be preferable to an “absolute fairness” that entails frequent and potentially dramatic changes, such as sudden progressions and regressions in their living standard. In other words, a world that is experienced as volatile may be less desirable than a certain degree of unfairness.

Yet, having algorithmic controls at their fingertips, individuals participating in a group may feel that they have no choice but to engage in a process of continuous negotiation and adaptation to rule-sets and social norms. The system’s affordances would engender an open cycle of societal self-adaptations and frequent change, inducing societal stress and fatigue. Nonetheless, since *efficiency* (i.e., speed) is a defining characteristic of ICT systems, an ICT-based solution could end up introducing additional and potentially thornier problems.

There are other important questions to address:

- How vulnerable would such ICT system be to ‘hijacking’ by external parties and what could be the consequences?
- When is fairness preferable to a certain degree of competition and could the system be re-configured to support either approach?
- Is the majority’s opinion always in the community’s best interest?
- Are there any collateral *costs* that such system would place on society?

Conclusions

Such questions and the ensuing design requirements must be carefully considered before irreversibly embedding societal governance in algorithmic technical systems. Since all possible scenarios cannot be predicted and addressed in advance, the ICT system itself must be sufficiently flexible to enable its evolution in parallel to the society it serves. If we can address such challenges, the potential rewards in empowering grassroots solutions to local issues (e.g., quality of experience in one’s living space) and coordinating collective action on a planetary scale (e.g., ensuring resource sustainability), are incalculable. But even then, given the dismal, unresponsive performance of the alternatives to algorithmic governance and self-organization, one could even simply ask: Can we afford not to?

References

- [1] Andrew Jones, Alexander Artikis and Jeremy Pitt, “The Design of Intelligent Socio-Technical Systems”. *Artificial Intelligence Review* 39(1):5-20, 2013.
- [2] Elinor Ostrom, *Governing the Commons: The Evolution of Institutions for Collective Action*, Cambridge University Press, 1990.
- [3] Michael Cox, Gwen Arnold and Sergio Villamayor Tomás, “A Review of Design Principles for Community-based Natural Resource Management”. *Ecology and Society* 15(4):1-38, 2010.
- [4] Jeremy Pitt, Julia Schaumeier, Alexander Artikis, “Axiomatization of Socio-Economic Principles for Self-Organizing Institutions: Concepts, Experiments and Challenges. *Trans. Auton. Adapt. Sys.* 7(4):1-39, 2012.
- [5] Jeremy Pitt, Didac Busquets and Sam Macbeth, “Self-Organising Common-Pool Resource Allocation and Principles of Distributive Justice,” *submitted*, 2014.
- [6] Nicholas Rescher, *Distributive Justice* (Bobbs-Merrill Publishing, 1966).
- [7] Lyria Bennett Moses, “Recurring Dilemmas: The Law’s Race to Keep up with Technological Change,” *Journal of Law, Technology and Privacy*, 2007(2):239-285, 2007.
- [8] Royal Commission on Environmental Protection (Chairman: John Lawton). *28th Report: Adapting Institutions to Climate Change*. The Stationery Office Limited, 2010.

Jeremy Pitt is a Reader in Intelligent Systems in the Department of Electrical & Electronic Engineering at Imperial College London, UK. His research interests are in

self-organizing multi-agent systems and their application to computational sustainability. He has collaborated widely, having worked on over 30 national and international projects, and has been involved in much inter-disciplinary research, having worked with lawyers, philosophers, psychologists, physiologists and fashion designers. He also has a strong interest in the social implications of technology, and had edited two volumes in this concern: This Pervasive Day (IC Press: 2012) and The Computer After Me (IC Press: 2014).

***Ada Diaconescu** is an assistant professor in the Computing and Networks department of Telecom ParisTech, in Paris, France. She is also a member of the CNRS LTCI research laboratory (Information Processing and Communication Laboratory). Her research interests include autonomic and organic computing, software engineering for self-adaptive and self-organising systems, component and service-oriented architectures, and interdisciplinary solutions for managing the complexity of cyber-physical systems. She received her PhD in computer science and electronic engineering from Dublin City University in 2006. Before joining Telecom ParisTech in 2009, she carried out various research projects at University of Grenoble, Orange Labs and INRIA Rhone Alpes.*