Algorithmic Self-Governance and the Design of Socio-Technical Systems

Jeremy Pitt[1], Dídac Busquets[1], Ada Diaconescu[2], Andrzej Nowak[3], Agnieszka Rychwalska[3], and Magda Roszczyńska-Kurasińska[3]

[1] Imperial College London, Exhibition Road, SW7 2BT, UK
[2] Telecom ParisTech, 46 rue Barrault F-75634 Paris Cedex 13, France
[3] University of Warsaw, ul. Stawki 5/7, 00-183 Warsaw, Poland

Abstract. The Digital Society is increasingly characterised by an ecosystem of smart, socio-technical applications. Unlike biological ecosystems, each application, and indeed the entire socio-technical ecosystem, is critically dependent on human-centred, mutually agreed, conventional rules for its effective and efficient operation, and inter-operation. This paper is concerned with exploring how to represent, reason with, and exploit these rules. In particular, it proposes the idea of *algorithmic self-governance*, which interleaves dynamic social psychology, holonic systems and self-organising electronic institutions, can provide a basis for developing socio-technical (eco)systems which empower solutions to large-scale collective action problems. We conclude by suggesting that this provides an innovative approach to the development of smart(er) cities.

Keywords: Socio-Technical Systems, Self-Organising Systems, Computational Social Intelligence, Electronic Institutions

1 Introduction

The Digital Society is increasingly characterised by an ecosystem of smart, socio-technical applications. These applications are predicated on the interaction of people and technology, and embedded in environments that are fully instrumented with devices and sensors, inter-connected (e.g. through both social and sensor networks) and intelligent (interleaving both social (human) and computational intelligence. Examples include electricity generation, distribution and storage, water management, and urban transportation, amongst others. The unification of these individual examples as an 'ecosystem' is well exemplified by the concept of *smart cities*.

Unlike biological ecosystems, each application, and indeed the entire socio-technical ecosystem, is critically dependent on human-centred, mutually agreed, conventional rules for its effective and efficient operation, and inter-operation. There is a well-established understanding of the importance of such conventional rules in the conduct of human affairs, especially when encapsulated by institutions. This understanding is, perhaps, best epitomized by the pioneering work of Nobel Laureate Elinor Ostrom [16], who identified an institution as a structured rule-set intended to regulate and/or constrain the behaviour of people, especially with regard to solving collective action problems like the long-term sustainability of a common-pool resource.

It has been a major challenge for computational social intelligence to understand, explain and engineer the processes underlying the formation, selection and modification of conventional rules for use in *electronic* institutions. The challenge ahead is leveraging this research in the convergence of computational intelligence with human intelligence in the representation of, and reasoning with, such rules in socio-technical systems. These systems would be especially beneficial in the resolution of collective action problems – for example, using local knowledge and behaviour to avoid undesirable macro-level outcomes and achieve desirable ones.

This paper is concerned with exploring how *algorithmic self-governance*, which interleaves dynamic social psychology, holonic systems and self-organising electronic institutions, as a basis for developing such socio-technical systems which empower local solutions to collective action problems. Section 2 considers the background and motivation to this work, including a critical analysis of Ostrom's work and its suitability for designing socio-technical systems. Section 3 surveys the three research areas contributing to the idea of algorithmic self-governance. Section 4 describes a case study in shared living spaces from which we derive our innovative proposal for developing smart(er) cities founded on, but going beyond, Ostrom's principles, which we call *Ostromopolis.*

2 Background and Motivation

Ostrom's pioneering work [16] showed how self-governing institutions could overcome the 'tragedy of the commons', which claimed to show that a group of appropriators with common, unrestricted access to a shared resource would inevitably act so as to deplete the resource in the short term, even if it was in no-one's interest in the long term. Based on extensive fieldwork, she showed how institutions (identified as structured rulesets which prescribe who could perform what actions in a specific 'decision arena' or 'action situation', what actions were permitted, proscribed or obliged, membership conditions, sanctions for not complying with the rules, etc.) could promote sustainability of a common-pool resource, without resorting to privatisation or centralisation.

Observing that the presence of a ruleset was not in itself a sufficient condition for enduring resource management, Ostrom identified common features of institutions which differentiated success stories from failures (for example, [16, p. 180], no clear membership boundaries, no support for self-determination, inadequate monitoring, or no support for 'efficient' conflict resolution). She then turned her attention to the problem of 'supply': faced with a common-pool resource management problem, there was no need to 'hope' that an institution with the requisite features for sustainable management would evolve. Instead, supported by an appropriate framework and accompanying tools and methods, institutions could be *designed* with these features specified as requirements.

Although an institution was supposed to identify who could perform what action in a specific 'action situation', Ostrom's work did not explicitly distinguish between physical capability, institutionalised power and permission (commonly made in the study of social, legal and organisational systems). However, by invoking the concept of *institutionalised power* [4], the design principles could be formalised in computational logic and used as an executable specification for *electronic* institutions for managing resource allocation in open computer systems and networks [19].

Moreover, the concept of *fairness* was more or less implicit in the operational choice rules for resource allocation – the relevant design principle only prescribed that those affected by these rules should participate in their selection, and assumed that those participating would presumably select rules that were, somehow, fair. For electronic institutions, the formalisation of Ostrom's principles was complemented by the formalisation of a theory of distributive justice [21] to ensure fairness in the distribution of resources [17].

In general, though, it could be argued that Ostrom's commitment to specifying institutions in concrete form, e.g. through principles, design methods and grammars, was rooted in political and economic science, but less so in computational, psychological and complexity sciences. As a result, her definition and analysis of 'action situations' overlooked not just fundamental organisational concepts such as institutionalised power, but also overlooked both the dynamic socio-psychological processes involved in the (bottom-up) emergence, (top-down) supply and (middle-out) self-adaptation of institutions [10, 9, 11], and the role of social networks in influencing decision-making in such situations [18].

It might also be argued that the design principles are well-suited to local situations, but not for situations that have multiple, deeply entangled priorities driven by possibly competing or even contradictory policy objectives, or when there are external authorities whose policies and policy demands have to be observed. However, Ostrom contended that large-scale collective action problems, with correspondingly large-scale outcomes, are not necessarily better addressed by top-down policy-making [14]. It was proposed that policies made at national and international level required local and regional action and enforcement, and governance should therefore be polycentric – i.e. composed of multiple centres of decision-making [13]. However, a comprehensive explanation of how polycentric governance can be identified, designed and delivered is still missing.

Finally, the interaction between computational intelligence and social intelligence (and technology in general) is also absent from Ostrom's original work. Given the criticality of the interface between users and their infrastructure [7], if that infrastructure is highly instrumented, as is the case in smart cities, then the human-computer interaction and ergonomics issues must also be considered.

It is (some of) these lacunae that we address in this work. In doing so, we aim to convert what might otherwise have been 'failures' into success stories, i.e. by designing and developing complex socio-technical systems with diverse computational and social intelligences for empowering successful collective action, using adaptive institutions that build on, but go beyond, Ostrom's principles.

3 Building Blocks for Algorithmic Self-Governance

From an abstract perspective, the objective of many socio-technical applications in the Digital Society (e.g. for infrastructure management, shared living spaces, and urban transportation) can be construed as managing a collective action situation. In general, a collective action situation involves several key features:

- it involves a group of people working together in a common space, but ...
- ... individuals may have a self-interest which conflicts with the group interest, which encourages free riding, and ...
- ... the costs of an action may fall on an individual, but the benefits accrue to the group, often requiring other incentives to contribute, for example in the form of social capital [15].

Starting from Ostrom's design principles for enduring institutions, we propose interleaving three building blocks for the design and development of socio-technical systems empowering successful collective action solutions: dynamic social psychology, electronic institutions, and holonic system architectures.

3.1 Dynamic Social Psychology

Dynamic Social Psychology is concerned with how dynamical systems, in which sets of components interact in complex, non-linear fashion but nevertheless produce coherent patterns, can be applied to social psychology, and has led to a number of theories concerning social change and social cognition:

- The Dynamic Theory of Social Impact, which specifies the processes by which a collection of private attitudes and beliefs becomes public opinion, common knowledge, or a form of culture [10].
- The Bubble Theory of Social Change, which specifies how a sustainable social change may be achieved, and concentrates on changing fragments of social networks (clusters or bubbles) rather than separate individuals [9]. In particular, Bubble Theory can be used to understand better the interaction between these structures.
- The Dynamic Theory of Societal Transition, defining the processes and conditions under which (meso-level) social structures are changed [11]. In particular a formal model of this theory will identify and specify how grassroots activists can control these processes in developing meso-level structures (i.e. institutions) that regulate or constrain micro-level behaviours to achieve desirable outcomes (and/or avoid undesirable ones).

To illustrate the principles and potential of dynamic social psychology for providing the theoretical foundations of designing socio-technical systems, Project ROSE (Regional Centres of E-learning) represents an example of an early attempt of programmed emergence of organisation [12]. The challenge was to promote the use of ICT, especially the Internet, in education in Poland. However, the rapid advances of ICT usually render any non-evolving educational program obsolete in just a few years. The solution was to create a learning community in the form of an expanding network of teachers that constantly adapted to new developments in ICT.

ROSE was based on the idea that teacher enhancement is a social change process rather than a transfer of knowledge. The Bubble Theory of Social Change [9] specifies how a sustainable social change may be achieved – by concentrating on changing fragments of social networks (clusters or bubbles) rather than separate individuals. ROSE was therefore a mixture of face-to-face workshops and Internet mediated interactions. The workshops enabled the teachers to learn to collaborate with each other and to develop trust. From each workshop several individuals were selected as natural leaders to seed the ROSE network. After the initial workshop the training was conducted over the Internet using an e-learning platform. The communication structure resembled a star with the university performing the role of the central hub, and each school being a spoke.

The leaders in each school initially worked with teachers from their own school but in the next stage schools already in ROSE collaborated with each other in the preparation of programmes for other schools. Meso-level structures (formal groupings with rules, roles, processes, designated groups responsible for decisions in specific areas; and informal groupings based on friendship circles, interest groups, and so on) emerged as clusters of collaborating schools, local administration and businesses etc. Afterwards, the meso-level structures grew stronger and bigger as more common initiatives were undertaken. The role of the university decreased as the network became increasingly decentralized.

In summary, project ROSE has exemplified the necessary conditions for *planned emergence*, namely multi-functional micro-level components (i.e. people able to fulfil different roles in different contexts); the formation, operation and dissolution of interacting meso-level structures (i.e. institutions); and the 'shaping' of the meso-level structures through which objectives at the macro-level can be achieved by collective, purposeful action at the micro-level. The open question is how to deliver planned emergence in socio-technical systems, which includes both computational and social intelligence as micro-level components, electronic institutions amongst the meso-level structures, and meso-level objectives which are global in nature (e.g. climate change). We begin to address this question by building on the concept of electronic institutions.

3.2 Electronic Institutions

Electronic institutions are used to represent the structures, functions and processes of an institution in mathematical, logical and computational form.

In terms of functional representation, an institution's rules can be divided into three levels, from lower to higher [16]: *operational-choice rules* (OC) are concerned with the provision and appropriation of resources, as well as with membership, monitoring and enforcement; *social collective-choice rules* (SC) drive policy making and selection of operational-choice rules; and *constitutional-choice rules* (CC) deal with eligibility and formulation of the collective-choice rules.

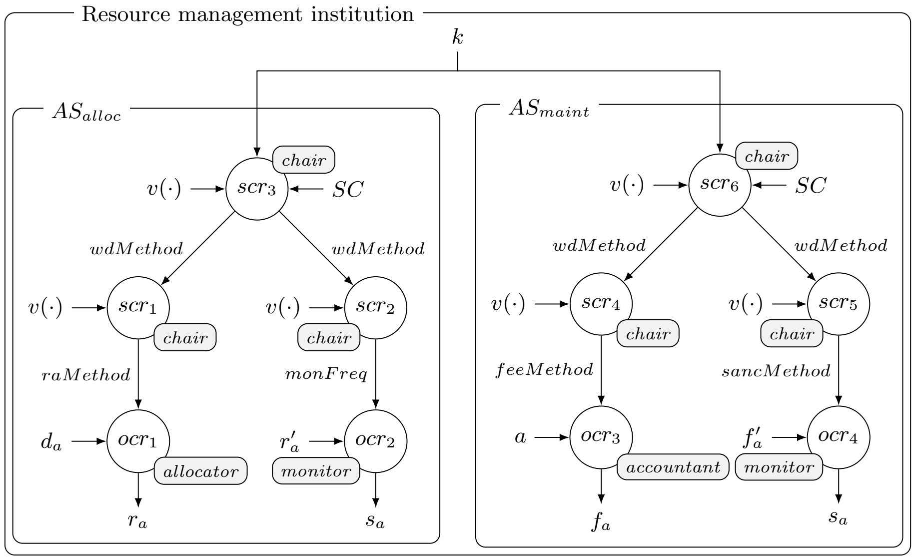

Fig. 1: Rules relationships: nodes represent rules, edges their inputs/outputs, and shaded rectangles the role empowered to execute the function.

For example, Figure 1 illustrates a resource management institution with two action situations, one for *resource allocation* (AS_{alloc}) and one for *infrastructure maintenance* (AS_{maint}). In AS_{alloc}, there are two operational-choice rules: ocr_1 *allocates* the resource to the users, according to their demands and some allocation method ($raMethod$); ocr_2 applies *monitoring* to identify any users that appropriate more resources than they have been allocated. For the social collective-choice rules, scr_1 selects the allocation method ($raMethod$); scr_2 selects the monitoring frequency; and scr_3 selects the *winner determination* method to be used in the voting procedures of scr_1 and scr_2. (The functions are similar for the infrastructure maintenance action situation AS_{maint}.)

A formal representation of institutional processes can also be given, which identifies their procedural, temporal and normative aspects, typically of concern in the study of social and organisational systems. In [19, 17], computational logic was used to represent these processes, using the Event Calculus (EC) [6]. This is fundamental to the representation of *self-organisation* and *self-governance*. However, a key challenge now is to encapsulate formal models of social processes, as specified by Dynamic Social Psychology [10, 9, 11], within the framework.

3.3 Holonic System Architectures

In terms of *engineering* algorithmic institutions for ‘real world’ socio-technical applications, we advocate the use of holonic system architectures. Holonic architectures and their key role in creating viable complex systems were introduced by Simon [23], refined by Koestler [5], and progressively adopted in software system engineering. For instance, holonic principles have been referred to as the "laws of artificial systems. Ignoring these laws is analogous to ignoring gravity in civil engineering" [24].

In brief, a holonic system is composed of simpler subsystems, which are composed of sub-subsystems and so on, recursively. Each system resource is 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. It also improves reactivity, stability and robustness by enabling local self-* processes and limiting their global effects. For example, a smart house 'system' at one level (i.e. a house with a smart meter installed and programmable devices) becomes a sub-system itself at the next scale up (e.g. a district with smart houses and other forms of renewable generation), while districts themselves are sub-systems at the next higher scale, and subject to a different set of policies and policy goals (see Figure 2).

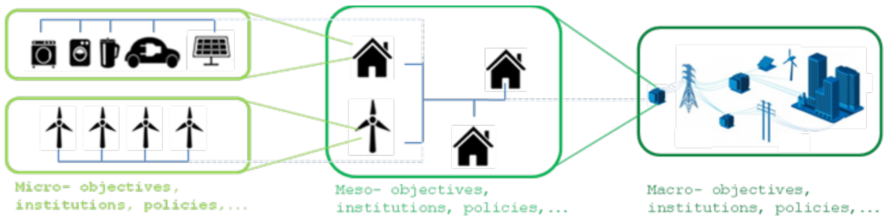

Fig. 2: Holonic System Architecture (e.g. for SmartGrid)

A holonic approach is required to address critical complex system issues, such as scalability, elasticity, adaptability, robustness, resilience and support for multi-scale, multi-objective policies, via recursive coordination of micro and macro processes. Furthermore, The holonic systems perspective provides an appropriate engineering paradigm not just for realising electronic institutions and planned emergence, but also in representing polycentric governance [13] and in dealing with psychological processes with span institutional boundaries. In addition, it would allow this kind of system and its inherent benefits to scale with the number and dynamicity of participants, which would make it applicable to smart city eco-systems and provide a better opportunity for the formation and development of social intelligence in contemporary socio-technical environments.

4 From Shared Spaces to "Ostromopolis"

This section briefly presents a case study in managing a shared living space as a common pool resource, where the design and development of a socio-technical system could benefit from the application of dynamic social psychology, electronic institutions, and holonic system architectures outlined in the previous section. From there, we briefly consider how to scale up from small local situations to socio-technical (eco)systems for larger contexts such as smart(er) cities.

4.1 Shared Living Spaces

Any shared living space, such as a communal flat, an open-plan office, or even a public space such as a park, require people to share a common space, where violation of (implicitly or explicitly stated) conventional rules, or social norms, can cause instances of incivility [20]. Such incivility, characterised by a low-intensity form of deviance from accepted norms, can be difficult to detect and resolve, but is also very harmful for the people who experience it regularly.

Therefore, it is a pressing problem in both ergonomics and urban planning to reduce the negative side-effects of incivility. The technological solution we have proposed for addressing the incivility problem, is MACS (M—s Affective Conditioning System): a system that attempts to avoid, reduce and/or resolve incivility before it escalates into a higher-intensity situation, e.g. conflict or aggression [22]. MACS is intended to emphasise stakeholder engagement and empower collective choice: firstly by avoiding micro-management, as incivility episodes are resolved between stakeholders (i.e. the occupants of the shared space themselves), and only as a last resort by appeal to higher authorities; and secondly by providing social support, through a network of communication and mutual obligations, via the collective selection, monitoring and enforcement of the stakeholders' own social norms and pro-social processes such as forgiveness [25].

We envision the shared living space as a *common pool resource* which we seek to manage according to the institutional design principles of Elinor Ostrom [16]. In this respect, the metaphor we are pursuing is that the (intangible) 'office ambience' is a pooled resource which the office occupants can deplete by anti-social behaviour and re-provision by pro-social behaviour. Furthermore, what is (and is not) anti-social behaviour is determined by the occupants themselves – a specific instantiation of Ostrom's third principle (that those affected by collective choice arrangements participate in their selection). Consequently, MACS implements a voting system for social norms, which allows for those (and only those) admitted to a shared space to vote positively or negatively for a norm. It also allows people to suggest new norms, as the dynamic nature of offices might mean there is a constant need to change norms, so MACS provides support for this process.

Figure 3(a) depicts the first screen displayed for a user, after a successful login to MACS. The navigation bar, on top, and the footer bar, at the bottom of the screen, are constant throughout MACS. The navigation bar provides direct access to the home screen, the social norms screen and the historical information about events where the logged-in user has been involved in, as an offender. Below the navigation bar, is the set of avatars representing all the people the logged used shares the workplace with. By hovering on each of the avatars the text "Flag *person's name*'s violation of norms" shows up, where *person's name* is replaced by the chosen person's name. By clicking on an avatar, the user is taken to the flagging screen, where they can create a new event, by flagging a violation of norms by the person they chose. At the bottom left area of the screen there are two different items regarding the logged user: their current reputation (standing with the community for compliance with norms) and its evolution graph for the previous 10 days, their avatar and their name.

A core function of MACS is to keep the users informed about the social norms they must abide by. Besides being able to check the norms at all times, users must also be able to vote for them, positively or negatively, and to suggest new norms. Figure 3(b) displays the "Social Norms" screen for an open plan office. Here all norms are presented, ordered by severity level, from the most to the least critical. Each norm is printed in the colour code that reflects its severity. Red means the norm is very critical, orangey-red means critical, orange means average, and finally yellow means minor. In this case, there aren't any minor severity norms to be displayed. In front of each norm, in square brackets, is its category. Categories are "noise", "privacy", "food", "environment", "politeness" and "borrowing items". Below each norm is its description. And finally by each norm are an approve (thumbs up) and a disapprove (thumbs down) buttons, which can be used to vote positively, or negatively, respectively, for the norm. At the bottom of the list of norms is the suggestion box, where the user may suggest a new norm for their workplace.

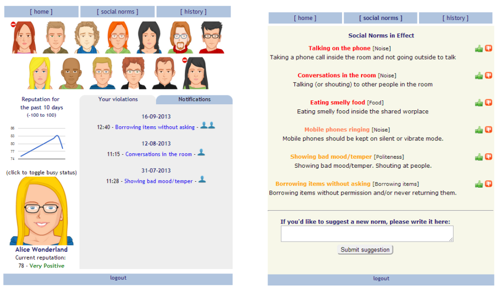

(a) (b)

Fig. 3: (a) MACS user start screen; (b) MACS social norms interface

To ensure that MACS meets its objective of reducing incidents of incivility in shared spaces, the contribution of formal models of social processes (e.g. conflict and forgiveness), self-governance by self-selection and modification of rules, and the requirement for a holonic systems approach (i.e. a user is a holon in a flat; a flat is a holon in a building; a building is a holon in a district, and so on) are all evident.

4.2 "Ostromopolis"

The aim of these case studies has been to show how the varied and cross-cutting 'building blocks' of algorithmic self-governance provide a foundation for an innovative approach to the design and development of socio-technical (eco)systems.

The current "big data" approach to making 'sense' out of vast amounts of conflicting and unstructured data flowing in from ICT devices deployed in various areas of social life is leading to some advances in predictive, and prescriptive analytics, with outcomes that range, arguably, from the beneficial and insightful to the unwarranted, alarming and intrusive.

The foundational and original character of the approach outlined here is to redefine the problem: instead of thinking of "global" as vast, unstructured and/or conflicting we define "global" as complex and holonic. The fundamental solution we offer is to empower the social structure at different levels of organization so that the self-organizing institutions may collaborate with policy makers to govern the smart infrastructures and the data they are generating through tailored ICT platforms implementing the electronic institution engine. This way, the "global" is deconstructed into local aggregates that themselves analyse, structure, interpret and utilise their data flows. In such a holonic architecture of global systems policies turn from single-paths set towards globally defined goals into constraints defining the boundaries for micro-governance at each level of the social structure.

In this way we believe it can be possible to develop socio-technical applications which empower users, and an ecosystem which unites these applications in managing multiple resources for the common good. We propose to start from, but go beyond Ostrom's theories, to overcome the limitations outlined in Section 2, to provide the foundations for promoting awareness, responsiveness and pro-social incentives for collective action in a socio-technical ecosystem for smart(er) cities. This innovative vision we call *Ostromopolis*.

5 Related Work

In realising the vision of Ostromopolis, there are, in fact, several other pieces of the 'jigsaw' required, beyond the 'building blocks' of Section 3. This includes:

- the social computer [8]: in which the designers of socio-technical applications seek to synthesise the intelligence of human and automated computational units to tackle so-called 'wicked' problems;
- social capital: the role of social capital [15] and the rise of cryptocurrencies such as Bitcoin and Venn, in the creation of incentives and alternative market arrangements has yet to be fully explored;
- serious games and gamification: gamification is a natural extension of serious games from artificial settings with self-contained game-defined rewards and "win" conditions, to real-life situations where the rewards and win conditions may be rather different. In real-life scenarios concerning common-pool resources, the "win" condition is very often sustainability, rather than termination of the game, i.e. the aim is to keep the game going.
- knowledge commons: Ostrom's design principles reflect a pre-World Wide Web era of scholarship and content creation, and despite some insightful work [3], these developments make it difficult to apply the principles to non-physical shared sources such as data or knowledge commons, and a

further extension of the theory is required to develop applications based on participatory sensing;
- privacy: new platforms which respect data privacy as a fundamental design principle are required, such as Open Mustard Seed (OMS) being developed by ID3 (The Institute for Data Driven Design) [2];
- scale: the scale of the applications requires systems to process many thousands of events per second. This is beyond the capacity of simple versions of the Event Calculus, and new dialect is required, such as the Run-Time Event Calculus (RTEC) [1].

6 Summary and Conclusions

In this paper, we started from Ostrom's pioneering work on self-governing institutions, but noticed there were, perhaps, some aspects where under-specified (or even not at all): for example the representation of empowerment, fairness, psychological processes, and polycentric governance.

The ROSE project has demonstrated that by applying principles of dynamic social psychology, people can be empowered to develop systems-of-systems, based on organisations and institutions, from the middle-out, which are user-centric and 'fit for purpose' because they are self-designed by and for the users. Based on this, we proposed that algorithmic models of such processes could be accommodated within the formal specification of self-organising electronic institutions, and furthermore, that a polycentric governance model for a 'system of systems' (of such institutions) could be realised using holonic system architectures.

In conclusion, it has been the aim of this position statement to indicate the opportunities, challenges and potential benefits of the cross-collaboration between the three research fields of dynamic social psychology, electronic institutions, and holonic systems. It is our contention that this original interdisciplinary composition can provide the foundations for designing and developing an (eco)system of socio-technical applications for smart(er) cities. This innovative proposal, i.e. founding smart cities on Ostrom's principles for self-governance and successful collective action, is what we have called *Ostromopolis*.

Acknowledgements

This work was supported by funds from Polish National Science Centre (project no. DEC-2011/02/A/HS6/00231). We would also like to thank the anonymous reviewers for their helpful comments.

References

1. Artikis, A., Sergot, M., Paliouras, G.: Run-time composite event recognition. In: Distributed Event-Based Systems (DEBS). pp. 69–80 (2012)
2. Hardjono, T., Deegan, P., Clippinger, J.: Social use cases for the ID3 open mustard seed platform. Technology & Society Magazine (2014)

3. Hess, C., Ostrom, E.: Understanding Knowledge as a Commons. MIT Press (2006)
4. Jones, A., Sergot, M.: A formal characterisation of institutionalised power. Journal of the IGPL 4(3), 427–443 (1996)
5. Koestler, A.: The Ghost in the Machine. Hutchinson Publisher (1967)
6. Kowalski, R., Sergot, M.: A logic-based calculus of events. New Generation Computing 4, 67–95 (1986)
7. Lam, W.F.: Governing Irrigation Systems in Nepal: Institutions, Infrastructure and Collective Action. Oakland, CA: ICS Press (1998)
8. Miorandi, D., Maggi, L.: “Programming” social collective intelligence. Technology & Society Magazine 33(3), 55–61 (2014)
9. Nowak, A., Lewenstein, M., Szamrej, J.: Bable modelem przemian spolecznych (bubbles: a model of social transition). Swiat Nauki (Scientific American Polish Edition) 12 (1993)
10. Nowak, A., Szamrej, J., Latane, B.: From private attitude to public opinion: a dynamic theory of social impact. Psychological Review 97, 362–376 (1990)
11. Nowak, A., Vallacher, R., Kus, M., Urbaniak, J.: The dynamics of societal transition: modeling non-linear change in the Polish economic system. International Journal of Sociology 35, 65–68 (2005)
12. Nowak, A., Winkowska-Nowak, K., Rycielska, L. (eds.): Szkoła w dobie internetu (Education in the age of Internet). Warsaw: PWN (2009)
13. 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. Stockholm: Nobel Foundation (2010)
14. Ostrom, E.: Thinking about climate change as a commons. In: 15th Annual Philip Gamble Memorial Lecture, pp. 1–34. UMass Amherst (2011)
15. Ostrom, E., Ahn, T.: Foundations of Social Capital. An Elgar Reference Collection, Edward Elgar Pub. (2003)
16. Ostrom, E.: Governing the commons: The evolution of institutions for collective action. Cambridge, UK: Cambridge University Press (1990)
17. Pitt, J., Busquets, D., Macbeth, S.: Distributive justice for self-organised common-pool resource management. ACM Trans. Auton. Adapt. Syst. 9(3), 14 (2014)
18. Pitt, J., Ramirez-Cano, D., Draief, M., Artikis, A.: Interleaving multi-agent systems and social networks for organized adaptation. CMOT 17(4), 344–378 (2011)
19. 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 (Dec 2012)
20. Porath, C., Pearson, C.: The price of incivility. Harvard Business Review 91(1-2), 114 (2013)
21. Rescher, N.: Distributive Justice. Bobbs-Merrill (1966)
22. Santos, M., Pitt, J.: Emotions and norms in shared spaces. In: Balke, T., Dignum, F., van Riemsdijk, M.B., Chopra, A. (eds.) COIN. LNCS, vol. 8386, pp. 157–176. Springer (2013)
23. Simon, H.: The architecture of complexity. Proc. American Philosophical Society 106(6), 467–482 (1962)
24. 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)
25. Vasalou, A., Hopfensitz, A., Pitt, J.: In praise of forgiveness: Ways for repairing trust breakdowns in one-off online interactions. Int. J. Hum.-Comput. Stud. 66(6), 466–480 (2008)