L&PS – Logic & Philosophy of Science
Vol. IX, No. 1, 2011, pp. 1-15

Modelling Resource Allocation in Open Embedded Systems: Some Logical and Epistemological Issues

Jeremy Pitt
Department of Electrical & Electronic Engineering
Imperial College London (UK)
j.pitt@imperial.ac.uk

Julia Schaumeier
Department of Electrical & Electronic Engineering
Imperial College London (UK)
j.schaumeier09@imperial.ac.uk

Marcello D'Agostino
Department of Economics and Management
University of Ferrara (Italy)
marcello.dagostino@unife.it

Published by L&PS – Logic and Philosophy of Science
http://www2.units.it/˜episteme/ – ISSN: 1826-1043

ABSTRACT. We consider the problem of deciding resource allocation policies in open embedded systems, where the system components have to form opportunistic alliances, share partial knowledge and collectively agree a policy that is congruent with the state of the environment in which the system is embedded. This problem arises in a number of systems and applications, including sensor networks, virtual organisations, cloud computing, swarm robotics, and infrastructure management. In this position paper, we outline a model for addressing this problem, and overview some of its features for which logical and epistemological issues concerning the dynamics of beliefs in a multi-agent system turn out to be relevant.

KEYWORDS: Resource Allocation, Logic, Epistemology

1. Introduction

This paper considers the problem of deciding resource allocation policies in open embedded systems, where the system components have to form opportunistic alliances, share partial knowledge and collectively agree a policy that is congruent with the state of the environment in which the system is embedded. Such systems are increasingly common for managing mobile ad hoc, sensor and vehicular networks [15], service-oriented systems such as virtual organisations and cloud computing [1], swarm robotics [16], and for delivering smarter demand-side infrastructure management, e.g. for water [2] or energy [18].

However, developing suitable models for this kind of situation (i.e. models which optimise 'fairness' and 'endurance' [12], which prevent strategic manipulation, or which promote security [11]) is a question for which logical, philosophical and epistemological research has particular relevance. In this position paper, we outline a model for addressing this problem and overview some of its features, for which logical and epistemological issues concerning the dynamics of beliefs in a multi-agent system provide deeper insight. In particular, we use key ideas from these research areas (scoring rules, affinity functions and depth-bounded reasoning) to map environmental states, partial knowledge and expressed preferences onto an actual resource allocation policy.

Accordingly, this paper is organised as follows. The next section presents a fuller description of the abstract problem specification, while Section 3 presents some elements of the proposed model. We conclude in Section 4 with some comments on how recent non-standard characterisations of classical logic could be employed to engineer more realistic models in which 'rational agents' are not assumed to be logically omniscient, but may be endowed with limited and non-uniform reasoning resources.

2. Abstract Problem Specification

There are many applications of open systems, including sensor networks, cloud computing, swarm robotics, and infrastructure management. In all these applications, the system components have to reason with partial information, and they have to share information and resources to achieve individual and collective goals. Moreover the system itself is embedded in an environment that is mutable by uncertain, exogenous events.

The operation of these systems can then be loosely divided into three stages: form an opportunistic alliance, decide a policy to provision and appropriate resources, and then perform the resource allocation, repeating as often as required. In previous work [12], we have advocated an *institutional* approach to this process, based on a logical axiomatisation of the socio-economic principles for common pool resource management [8]. One of these principles concerned *congruence*, specifically that the appropriation and provision rules should be congruent with the prevailing state of the environment in which the system is embedded. Since these are open, distributed and decentralised systems comprising autonomous, heterogenous and possibly competing components, this requires that they sense their local environment, communicate their beliefs to other nodes, and express a preference for the resource allocation policy that best suits the local environment. The collection of expressed preferences is then used to make a decision on an actual resource allocation policy.

However, the components themselves may perform actions which directly affect the environment, and the environment state may also change due to exogenous events beyond the nodes' control. Therefore, the operational resource allocation policy has to suit the current state of the environment, but the state of that environment is dependent upon the uncertain occurrence of significant (environment-changing) events. Accordingly, we require that the nodes make a forecast, as a subjective degree of belief, about the likely occurrence of these events. This introduces a 'security' problem, as it exposes the policy selection process to strategic manipulation: a node can be deliberately deceptive about the occurrence of an event to persuade others to decide on a policy that is beneficial to the deceiver but not in the common interest.

Minimally, this general (abstract) resource allocation problem can be formulated as a set of n nodes (agents) $\mathscr{A} = \{a_1, \ldots a_n\}$ forming k clusters (opportunistic alliances) $\mathscr{C} = \{C_1, \ldots C_k\}$, such that $\forall i, i \leq 1 \leq k, C_i \subseteq \mathscr{A}$. Picking one cluster $C \in \mathscr{C}$, we formally define a self-organising *institutional cluster* at time t as a 4-tuple:

$$\mathscr{I}_t = \langle C, \varepsilon, L, m \rangle_t$$

where:

- C is the set of agents;
- ε is the environment, a pair $\langle B, F \rangle$ where B is the set propositions whose truth-values are determined by the physical state, and F is the set propositions whose truth-values are determined by conventional agreement[1];
- L, is the resource allocation 'legislature', the set of rules by which agents are allocated resources; and
- m is a partial function $C \mapsto [0, P]$ which specifies the amount of resources allocated to each agent in C.

Assuming the system operates in time-slices, in each time-slice an agent may make a request for resources. The problem for the institutional cluster (or rather its members) is to compute the mapping m for the current time slice t based on the operational choice rules specified in L, and then select the operational choice rules for the next time-slice $t+1$ based on forecasts for the occurrence of uncertain events which may change the environment.

For example, the set of possible operational choice rules (policies) for the resource allocation could be:

[1] B and F are respectively called in philosophy the set of 'brute' and 'institutional' facts [17], with the institutional facts being the product of exercising *institutional power* [6]; i.e. the performance of a designated action in a specific context by an empowered agent, usually occupying a certain role.

- largest first: the agent making the largest resource request is allocated first, then the next largest, etc.;
- smallest first: the agent making the smallest resource request is allocated first, then the next smallest, etc.;
- first-come-first-served: the requests are allocated in the order in which they are generated;
- in turn: each agent is allocated resources in rotation;
- priority: the request with the highest priority is allocated first;
- ration: each request is partially allocated.

The best policy is determined by the current amount of resources to be allocated P, and the replenishment rate of the resource from the environment. This is determined by exogenous events whose occurrence is uncertain. For the sake of example, let us suppose that the replenishment rate can be *low*, *normal* or *high* as a result of certain events (in the water distribution CPR systems studied by Ostrom [8], the events could be weather-related, storms, rains, sunshine, etc.).

In a centralised, closed or co-operative system, such as an operating system, this is a scheduling problem which can be solved by well-known algorithms. In a distributed, open and competitive system, where the components are autonomous, heterogeneous and may (by accident or design) not comply with the provision or appropriation rules, we require a different approach in keeping with the use of institutions. Our proposal is discussed in the next section.

3. Logical, Philosophical and Epistemological Issues

In this section, we examine some features of the model for which important logical and epistemological issues concerning the dynamics of beliefs in a multi-agent system are particularly relevant. This includes scoring rules, affinity functions, depth-bounded reasoning, and judgement aggregation.

3.1. Scoring Rules

Let E_i, with $i \geq 1$, be a sequence of uncertain events. Essentially, a *scoring rule* is a function $f(a_i, e_i)$, where a_i is a real number in $[0,1]$ that represents a forecaster a's expressed opinion (subjective degree of belief) about the occurrence of E_i, and e_i is a Boolean variable such that $e_i = 1$ if E_i actually occurs and $e_i = 0$ if E_i does not occur. Then the forecaster is charged (or symmetrically rewarded) $f(a_i, 1)$ if E_i occurs and $f(a_i, 0)$ if E_i does not occur. Such rules are often employed in evaluating the accuracy of probabilistic forecasters and measuring their predictive success[2].

Applying this to the problem specified in the previous section, we might have three independent environment states, ε_0, ε_1 and ε_2, and events causing transitions between states, i.e. $E_{0,0}$, $E_{0,1}$, etc., as shown in Figure 1.

Then we require each agent, in the corresponding environmental state i at time-slice t, to express it's subjective belief in the likelihood of the events $E_{i,i}$, $E_{i,j}$ and $E_{i,k}$. These predictions are evaluated using the scoring rules in the following-time slice, and the agents rewarded accordingly. Note the reward might simply be an increase in 'reputation' and that this might effect implicit roles in the cluster such as leadership. Furthermore, if the scoring rule is proper, rational agents who aim to maximise their reward are forced to make forecasts that coincide with their true beliefs.

3.2. Opinion Formation

We now have a way of eliciting an opinion from an agent which precludes strategic manipulation, we also want to consider the influence of the opinion and the *social network* within the cluster.

---

[2] De Finetti showed that certain mathematical forms of the scoring rule prevent rational agents from 'cheating', that is, from expressing opinions about the occurrence of the uncertain event E that are different from their 'true beliefs' (on this point, see [7]). Scoring rules that have this property are called *proper*. If a non-proper scoring rule is used, there is room for strategic behaviour. But proper scoring rules can have various mathematical forms. De Finetti also showed that if one specific proper scoring rule is used (the quadratic scoring rule, called also Brier's rule) then subjective degrees of belief must obey Kolmogorov's axioms of probability, that is rational agents ought to reason in accordance with the laws of probability.
However, this does not happen if other proper scoring rules are adopted. On the other hand D'Agostino and Sinigaglia [5] showed that, under certain intuitive axioms, the quadratic scoring rule is uniquely determined as the only one that can be appropriately used in this context.

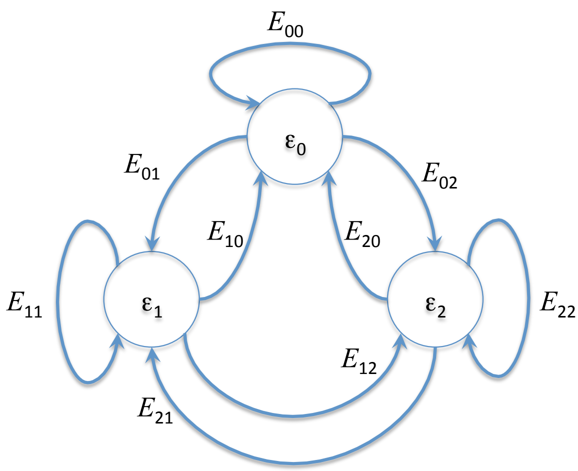

FIG. 1: *Events and Environment States*

In previous work, we have used use the opinion formation model of [14], whereby an agent i increases its confidence in another agent j based on how well j's opinion coincides with i's *mindset*. Assuming a positive evaluation for those opinions matching agent i's mindset and a negative for those contradicting it, then it can be said that the confidence in an exchange partner j increases as j's opinion matches i's mindset.

In this model, the opinion formation dynamics occurred at discrete time points and on a per issue basis. At each time point each agent exchanges opinions with other agents. An agent i's opinion changes at time $t+1$ by weighting each received opinion at time t with the confidence in the corresponding source (including its own opinion weighted by its self-confidence). Then in each time step, the *affinity* between agents can be different for each ordered pair of agents corresponding to the fitness between opinions and mindsets.

Furthermore, affinity affected confidence, so the confidence in other agents was redistributed accordingly. The 'distance' between one agent's mindset and another agent's opinion (as expressed to that agent) was used to update the affinity of the former agent for the latter. Confidence changed in time differently for each agent, based on the affinity between agents. Agents increased the confidence in those agents whose opinions fit their mindset.

In this context the work of [3] on distance functions appears to be relevant and allow for a natural extension of the model. In [14] affinity was supposed to compare an agent's mindset with another agent's expressed opinion concerning a single issue. This was a simplifying assumption of [14] , but it is natural to consider more realistic models in which affinity is evaluated with reference to a *set* of issues rather than a single one. In this approach, an agent's mindset or opinion on given set of n issues is expressed as a real-valued n-dimensional vector and affinity is therefore a function from $\mathbb{R}^n \times \mathbb{R}^n$ to $\mathbb{R}$. Here the background problem is the same as in the theory of multidimensional spatial voting, where voters are supposed to choose the candidates whose expressed opinions on the importance of a given set of issues have a minimal distance from the voter's 'true' opinions. D'Agostino and Dardanoni [3] have recently shown that, under some reasonable assumptions, the appropriate distance function must be some monotonic transformation of *Euclidean distance*, if all the issues are assumed to have the same 'relevance' for the agent who is making the affinity evaluation, or of *weighted Euclidean distance*, if some issues are more relevant than others (and the weights express their relative relevance). So, in the general case, affinity should be measured by a function f_a of the following form:

$$f_a = K\Big(\sum_{i=1}^{n}(w_{a.i}a_i - w_{a.i}b_i)^2\Big),$$

where: a is the agent who is making the affinity evaluation, $w_{a.i}$ is the weight assigned by a to issue i, a_i represents a's mindset about i, and b_i the expressed opinion of agent b about i.

3.3. Depth-bounded Reasoning and Belief Revision

Given that each agent in the institutional cluster expresses an opinion, and that opinion influences other agents, the next issue to address is how an agent may be prompted to change its opinions so as to trigger a global revision process that may lead to a change in the expressed preferences and, therefore, in the resource allocation policy.

An opinion, valued in $[0,1]$, is an agent's subjective degree of belief in the truth of a given proposition, i.eȧ fuzzy notion. However, agents often need to express their preferences by means of crisp propositions, which may in turn result from the application of inference rules whose premises and conclusions are also crisp. So we assume that each agent is equipped with a defuzzification method to turn their fuzzy opinions into crisp beliefs that are represented by *signed formulas*, i.e. expressions of the form $\mathbf{t}\,P$ or $\mathbf{f}\,P$, where P is arbitrarily complex formula built up from the atomic formulas of the language by means of the usual Boolean operators. The set of such crisp beliefs forms the agent's *belief database*. Moreover, we assume that agents are able, *in various degrees*, to extend their belief database and to detect inconsistencies by performing deductive inferences. Finally, we assume that each agent has the capability of applying additional inference rules that enables it to derive signed atomic atomic formulas from finite set of signed atomic formulas.

For example, suppose that, in a given agent's belief database, we have signed formulas of the form:

$$\mathbf{t}\,\mathit{common\_pool}(\mathit{high}), \mathbf{t}\,\mathit{common\_pool}(\mathit{normal}), \mathbf{t}\,\mathit{common\_pool}(\mathit{low})$$

which indicate that the agent has a (sufficiently strong) subjective belief that is is true that the common pool resource P is, in some sense, respectively 'high', 'normal' and 'low'. (There may also be integrity constraints required to ensure that the database does not contain incompatible crisp beliefs such as $\mathbf{t}\,\mathit{common\_pool}(\mathit{high})$ and $\mathbf{t}\,\mathit{common\_pool}(\mathit{low})$.) Depending on the observation of events, the opinion formation, and the defuzzification process, it will also have a subjective belief on the replenishment rate, i.e. $\mathbf{t}\,\mathit{rep\_rate}(\mathit{high})$. Then it would be able to infer an overall belief in the availability of resources by means of the following inference rule:

$$\frac{\begin{array}{l}\mathbf{t}\,\mathit{common\_pool}(\mathit{high})\\ \mathbf{t}\,\mathit{rep\_rate}(\mathit{high})\end{array}}{\mathbf{t}\,\mathit{resources}(\mathit{high})}$$

Or, suppose the agents measure their satisfaction which goes up or down according to whether or not they receive resources in a time-slice [13]. They also communicate this satisfaction as an opinion, and the agents can form a (subjective) overall assessment of whether the satisfaction in the cluster is high or low. There might be additional inference rules of the following form that allow agents to form a belief in the preferred operational choice rule (see next section):

$$\frac{\mathbf{t}\,resources(high) \quad \mathbf{t}\,satisfaction(high)}{\mathbf{t}\,ocr(largest\_first)} \qquad \frac{\mathbf{t}\,resources(high) \quad \mathbf{t}\,satisfaction(normal)}{\mathbf{t}\,ocr(smallest\_first)} \qquad \frac{\mathbf{t}\,resources(normal) \quad \mathbf{t}\,satisfaction(low)}{\mathbf{t}\,ocr(first\_come\_first\_served)}$$

$$\frac{\mathbf{t}\,resources(normal) \quad \mathbf{t}\,satisfaction(high)}{\mathbf{t}\,ocr(in\_turn)} \qquad \frac{\mathbf{t}\,resources(low) \quad \mathbf{t}\,satisfaction(normal)}{\mathbf{t}\,ocr(priority)} \qquad \frac{\mathbf{t}\,resources(low) \quad \mathbf{t}\,satisfaction(low)}{\mathbf{t}\,ocr(ration)}$$

Now, the problem is that we cannot realistically assume that each agent, at a given time slice t, 'believes' all the signed formulas that can be inferred from the set of formulas that are explicitly stored at t in the belief database, since the corresponding problem is computationally unfeasible. Moreover, for the same reason, an agent's database may be inconsistent, but the agent may be unable to detect this inconsistency and so find no motivation for revising its beliefs. A realistic model requires the possibility of *grading* the inferential power of agents, so that different agents may be able and/or willing to spend resources on performing deductive inferences up to different degrees of complexity. An agent that only reasons 'shallowly' may remain unaware of inconsistencies and so not revise its beliefs, whereas an agent that commits more of its own resources to reasoning 'deeply' might detect such inconsistencies and so start a belief revision process that may result in changing its expressed preference for the resource allocation policy (the operational choice rule).

A promising approach for solving this problem seems to be the one presented in [4], which allows for a natural proof-theoretical characterisation of an infinite hierarchy of tractable approximations to Boolean logic in terms of 'depth-bounded' natural deduction. This hierarchy of logical systems may be associated with a corresponding hierarchy of agents endowed with increasing inferential power. We maintain that replacing a complete deduction system for Boolean logic with a suitable depth-bounded approximation system and augmenting the latter with a suitable revision algorithm yields a logical model in which (i) agents can be realistically assumed to 'believe' all the depth-bounded logical consequences of their belief databases only up to a certain depth that is characteristic of each agent; (ii) when the database is inconsistent, belief revision may or may not be triggered depending on whether the inconsistency can be detected at the depth that bounds the inferential power of the agent.

3.4. Judgement Aggregation

After belief revision, the agents in a cluster will have to use their (new) beliefs to make judgments about logically interconnected propositions. However, a judgement on one proposition is not independent on judgements on some other propositions. Aggregating such judgements is not straightforward and leads to a variety of paradoxes (most notably the discursive paradox (the observation that majority voting applied to premises may yield a different outcome to majority voting on a conclusion).

The performance of different aggregation procedures for truth-tracking (such as the preference based procedure, the conclusion based procedure, and belief merging) has been investigated [9], and a formal characterisation of a voting protocol from the perspective of institutionalised power is given in [10]. Furthermore, there are many algorithms for single winner-determination given a list of candidates or alternatives. This is known as social choice or mechanism design: common alternatives include negotiation, auctions, and voting.

Five well-known methods based on voting are:

- *Plurality*. Each voter selects one candidate; the most named candidate wins.

- *Runoff*. Each voter selects one candidate in the first round. Top two (unless one already has a majority) enter a second round. Each voter selects one candidate (of two) in the second round; the most named candidate wins.

- *Borda*. Each voter rank orders all candidates. The Borda score for each candidate is computed from the accumulation of Borda points associated with each vote (with N candidates, rank k scores $N-k+1$ Borda points). The candidate with the highest Borda score wins.

- *Instant Runoff*. Each voter rank orders all candidates. The candidate with the least number of top-ranked vote is eliminated. The process is repeated until only one candidate remains.

- *Approval*. Each voter selects a subset of the candidates; the most named candidate wins.

It may be that each winner determination method (*wdm*) will return a different winner form the same set of votes. For example, given the following list of candidate votes for a set of policy choices $\{a,b,c,d,e,f\}$ (where a represents *largest_first*, b represents *smallest_first*, and so on):

$$[[a,b,c,e],[a,d,b,e],[a,c,d,e],[b,c,d,e],[b,d,c,e],[c,d,e,b],[d,c,e,b]]$$

we get the comparative results using a Prolog implementation, see Table 1.

method	winner	prolog inferences	robustness[19]
plurality	a	33	low
runoff	b	89	high
borda	c	281	low
instant (runoff)	d	328	high
approval	e	194	low

TABLE 1: *Comparison of Winner Determination Methods*

Thus it matters which winner determination method is used, not only because it may affect the selected winner, but the different methods have different resource requirements to compute, and also have a different level of robustness, as a measurement of its resistance to strategic manipulation [19].

4. Summary and Conclusions

In this position paper, we have presented a preliminary study of aligning resource allocation policies to the prevailing state of the environment, in resource-constrained open embedded systems. We considered a scenario in which a cluster (group) of agents have to (collectively) select rules for resource provision and appropriation, the method by which this selection is made, and (individually) decide how much of their resources to contribute to the performance of the job and the collective decision-making. We have suggested how suitable models of scoring rules, depth-bounded reasoning and judgement aggregation may fruitfully be combined to map environmental states, partial knowledge and expressed preferences onto an actual allocation.

There are several applications of logic in this scenario. Clearly, it plays a key role in the belief revision and in the judgement aggregation processes. However, the development of realistic models prompts for a departure from the mainstream approaches that can be found in the logic literature. In particular, we want to model 'population distributions' in which agents are assumed to have realistic and non-uniform reasoning capabilities, while it is widely agreed that logical systems are *idealisations* and, therefore, not intended to model the actual behaviour of rational agents.

This gap between the idealised assumptions of traditional logic and the observed inferential behaviour of real-life agents stands out as highly problematic in several areas of economics, sociology, psychology and political science. Granting that a certain amount of abstraction and simplification is necessarily involved in any scientific model, real applications require an approach to logical systems that (*i*) takes resource limitations and environmental constraints seriously and (*ii*) allows for indefinitely increasing *degrees of idealisation*. Such a graded approach would put researchers in a position to model situations where certain agents are capable of committing resources to reasoning 'deeply', and so bring to light non-trivial inconsistencies that trigger the belief revision process, compared to other agents which reason 'shallowly', and so may remain unaware of inconsistencies whenever these are deeper and therefore (computationally) harder to detect. We maintain that the revisitation of classical propositional logic put forward in [4], leading to a natural hierarchy of depth-bounded tractable logics of increasing complexity, can be sensibly employed to model a corresponding hierarchy of resource-bounded logical agents.

Moreover, as we try to lift from autonomous systems management to *self-aware* systems management, we conjecture that a proper treatment of depth-bounded reasoning in the style of [4] – given the link between awareness and tractability – may turn out to be useful to understand the proper manifestation of self-awareness in multi-agent systems.

REFERENCES

[1] ARDAGNA, D., PANICUCCI, B., and PASSACANTANDO, M. (2011), "A game theoretic formulation of the service provisioning problem in cloud systems", In *WWW '11*, pp. 177–186.

[2] BOULET, R., MAZZEGA, P., and JOUVE, B. (2009), "Environmental, social and normative networks in the MAELIA platform", In *Proc. ICAIL Worshop Legal & Negotiation Decision Support Systems*, M. Poblet, U. Schilds, and J. Zeleznikow, Eds., pp. 83–93.

[3] D'AGOSTINO, M. and DARDANONI, V. (2009), "What's so special about euclidean distance? A characterization with applications to mobility and spatial voting", *Social Choice and Welfare 33,* 2, pp. 211–233.

[4] D'AGOSTINO, M. and FLORIDI, L. (2009), "The enduring scandal of deduction: Is propositional logic really uninformative?", *Synthese 167,* 2, pp. 271–315.

[5] D'AGOSTINO, M. and SINIGAGLIA, C. (2010), "Epistemic accuracy and subjective probability", In *EPSA Epistemology and Methodology of Science*, M. Suárez, M. Dorato, and M. Rédei, Eds., pp. 95–105.

[6] JONES, A. and SERGOT, M. (1996), "A formal characterisation of institutionalised power", *Journal of the IGPL 4,* 3, pp. 427–443.

[7] MURA, A., DE FINETTI, B., GALAVOTTI, M., and HOSNI, H. (2008), *Philosophical Lectures on Probability*, Springer.

[8] OSTROM, E. (1990), *Governing the Commons*, CUP.

[9] PIGOZZI, G. and HARTMANN, S. (2007), "Aggregation in multi-agent systems and the problem of truth-tracking", In *Proceedings AAMAS '07*, pp. 674–676.

[10] PITT, J., KAMARA, L., SERGOT, M., and ARTIKIS, A. (2006), "Voting in multi-agent systems", *Computer Journal 49,* 2, pp. 156–170.

[11] PITT, J., RAMIREZ-CANO, D., DRAIEF, M., and ARTIKIS, A. (2011), "Interleaving multi-agent systems and social networks for organized adaptation", *Computational and Mathematical Organization Theory*, pp. 1–35.

[12] PITT, J., SCHAUMEIER, J., and ARTIKIS, A. (2011), "The axiomatisation of socio-economic principles for self-organising systems", In *Fifth*

IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO), pp. 138–147.

[13] PITT, J., SCHAUMEIER, J., AND ARTIKIS, A. (2011), "Role assignment in institutional clouds for rule-based enterprise management", In *Proceedings of the 5th International Conference on Rule-based Modeling and Computing on the Semantic Web, RuleML11*, pp. 237–251.

[14] RAMIREZ-CANO, D. and PITT, J. (2006), "Follow the Leader: Profiling Agents in an Opinion Formation Model of Dynamic Confidence and Individual Mind-Set", In *Intelligent Agent Technology (IAT '06)*, pp. 660–667.

[15] RAYA, M. and HUBAUX, J.-P. (2007), "Securing vehicular ad hoc networks", *Journal of Computer Security 15,* 1, pp. 39–68.

[16] SCHLACHTER, F., MEISTER, E., KERNBACH, S., and LEVI, P. (2008), "Evolve-ability of the robot platform in the symbrion project", In *Proceedings of the 2008 Second IEEE International Conference on Self-Adaptive and Self-Organizing Systems Workshops*, IEEE Computer Society, Washington, DC, pp. 144–149.

[17] SEARLE, J. (1969), *Speech Acts*, CUP.

[18] STRBAC, G. (2008), "Demand side management: Benefits and challenges", *Energy Policy 36,* 12, pp. 4419–4426.

[19] TIDEMAN, N. (2006), *Collective decisions and voting: the potential for public choice*, Ashgate Publishing Ltd.