

Micro-Social Systems: Interleaving Agents, Norms and Social Networks

Moez Draief and Jeremy Pitt and Daniel Ramirez-Cano¹

Abstract. Ad hoc networks can be formed from arbitrary collections of individual people (forming online computer-mediated communities), mobile routers (forming data communication networks) or electronic business processes (forming virtual enterprises).

One way to deal with common features of dynamism in the network topology and membership, conflicts, sub-ideal operation, security, and the general need for continuous operation in the absence of a centralised facility, is to treat the ad hoc network as a norm-governed multi-agent system and use participatory adaptation as the mechanism for achieving autonomic capability (i.e. a global system response derived from the collective local behaviours and interactions of the individuals comprising the system).

Therefore, complementing the formal representation of organisational behaviour defined in terms of roles, rules, norms, etc., this autonomic capability is at least partially derived from an underlying social network which plays a significant role in determining how, for example, conflicts are resolved and how the organisation itself is run.

This position statement presents initial developments in what we call *micro-social systems*, which arise from interleaving a logical model of norm-governed systems with a mathematical model of social networks, and its application to issues of resource allocation, security, conflict resolution and self-adaptation in ad hoc networks.

1 INTRODUCTION

Ad hoc networks can be formed from arbitrary collections of individual people (forming online computer-mediated communities), mobile routers (forming wireless data communication networks) or electronic business processes (forming virtual enterprises).

All three types of ad hoc network exhibit similar features which require attention, for example:

- **Dynamism:** the network topology and nodes can vary rapidly and unpredictably.
- **Conflicts:** the network consists of heterogeneous nodes which may be competing rather than co-operating. This might give rise to conflicts over opinions, priorities, contracts, and so on.
- **Sub-ideal operation:** the nodes themselves may fail to comply according to the system specification, by accident, necessity, or design.
- **Security:** a successful operation may be subject to attack, either from malice or for profit. Examples include deception in online communities, denial of service in ad hoc networks, untrustworthiness in virtual enterprises, and so on.
- **Continuity:** even if all the network nodes change, it may be desirable for the network itself to be recognisably and identifiably “the

same”, offering the same support, functionality, services, and so on.

However, we anticipate that there is likely to be no permanently-connected, centralised facility (e.g. a moderator in an online community, a network server, or a central database, file manager or other co-ordinating mechanism in a virtual enterprise) to resolve any of these issues; unsurprisingly, of course, since these are essentially peer-to-peer networks.

Various proposals to deal with these issues have been presented, which analyse the ad hoc network as an open system (see [4] for a survey). In the same vein, we have advocated the use of norm-governed multi-agent systems [5] for formal specification of open systems, and more recently to facilitate adaptation by voting to change the specification [8]. In particular, following ideas of Ostrom [25] we proposed to use *participatory adaptation* as the mechanism for achieving autonomic capability (i.e. a global system response derived from the collective local behaviours and interactions of the individuals comprising the system).

However, norm-governed system specification provided the framework for (provably) correct procedures for casting and counting votes [26], and declaring the winner, while mechanism design provided the algorithm(s) for mapping expressed preferences onto a collective choice (computational social choice [9]). This formal analysis, from an external perspective, needs to be complemented by data structures and protocols which enable individual agents to determine a preference to express, usually based on information exchanged with their peers. Therefore, complementing the formal representation of organisational behaviour defined in terms of roles, rules, norms, procedures, preferences, etc., the requisite autonomic capability is at least partially derived from an underlying social network which plays a significant role in determining how, for example, conflicts are resolved and how the organisation itself is run.

This position statement presents initial developments in what we call *micro-social systems*, which arise from interleaving a logical model of norm-governed systems with a mathematical model of social networks, and its application to issues of resource allocation, security, conflict resolution and self-adaptation in ad hoc networks. We begin in Section 2 with a review of relevant background, and continue in Section 3 with a discussion and definition of micro-social systems. We then discuss (Section 4) a sample scenario in security in ad hoc networks, paying attention to the problem that this is a resource-constrained environment (i.e. power is a limited resource). We then present a proposed adaptive network security scheme, before drawing some conclusions and outlining our future programme of research in Section 5, in particular considering the inter-play and interleaving of overt organizational structures and the underlying social networks.

¹ Department of Electrical & Electronic Engineering, Imperial College London, UK, SW7 2BT; email: {m.draief,j.pitt,d.ramirez}@imperial.ac.uk

2 BACKGROUND

There is a rich tradition of research in multi-agent systems which is both informing and motivating the current proposal. This includes norm-governed systems, socio-cognitive systems, and information flow in social networks.

2.1 Norm-Governed Systems

A norm-governed multi-agent system can be expressed in terms of a set of agents (the members of a society), a set of social constraints on a society (norms, and other constraints, such as physical and logical constraints), a set of roles that members can play, a communication language, relationships between the members, including power, ownership and representation relations, and the structure of the society, including hierarchies, sub-structures, and ontologies for domains of responsibility.

We maintain the standard and long established distinction (in legal, social, and organizational systems research) between physical capability, institutionalised power and permission. Accordingly, a specification of the social constraints of a norm-governed multi-agent system expresses four aspects of agent activity: (i) the physical capabilities; (ii) institutionalised powers; (iii) permissions, prohibitions and obligations of the agents; and (iv) the sanctions and enforcement policies that deal with the performance of prohibited actions and non-compliance with obligations.

On this basis, we have developed a framework for executable specification of open systems [5], and an accompanying suite of protocols for norm-governed interaction, including resource allocation [2], negotiation [5], and voting in deliberative assemblies for virtual organizations [26].

In the most recent work, we have addressed the issue of organized adaptation in norm-governed systems, and have developed a norm-governed meta-protocol for changing object-level protocols and parameters [3]. The approach has also been applied to develop legal processes for virtual organizations; of particular relevance to the current context were protocols for conflict prevention and alternative dispute resolution [27].

2.2 Socio-Cognitive Systems

A socio-cognitive multi-agent system is a system in which the social relationships between agents are conditioned by cognitive, economic, or emotive factors. For example, in [23] we defined a framework which accounted for socio-cognitive and socio-economic factors to inform a trust decision in e-commerce. This showed how ‘first encounter’ decisions were best dealt with by a complex calculation based on risk exposure and recommendations (‘risk’ trust); but how ‘*n*th encounter’ decisions were short-cut based on personal experiences (so-called ‘reliance’ trust). In [31] we combined both socio-cognitive and socio-emotive approaches to develop a formal model of forgiveness complementary to the trust model. Forgiveness turns out to be an essential element of autonomic agent systems as a repair mechanism for when a trust decision goes wrong: similarly, when a security breach occurs, the most effective response might not necessarily be to engage a ‘stronger’ security level.

We have also been concerned with applications of multi-agent systems which depend on a set of distributed nodes exchanging information to reach some collective decision [28]. In open systems, though, self-interested agents might resort to different strategies, such as ‘lying’, in order to shift the collective decision towards that which will

benefit their individual goals. We address this problem by defining such systems in terms of social exchange. Then, by making certain assumptions about the individual nodes and the collective decision-making, we define an opinion formation model of the system. We can then show how characteristics of the model can be matched with parameters of the system to achieve certain desirable properties of the information exchange, e.g. consensus, non-dominance, etc.

The study of norm-governed and socio-cognitive systems converged in the idea of participatory adaptation [7, 8]. This work used an iterated ‘tragedy of the commons’ scenario, with a dynamic network, partial knowledge, no central control, and self-interested agents. Two votes were taken: one for whom to allocate resources, and one to decide how many votes should be received in order to be allocated resources. The idea was that co-operative agents should manage the system by voting ‘fairly’. Initial experiments showed that ‘responsible’ agents performed better than selfish or cautious ones [7] and that social networking (gossiping) algorithms can be used on an individual and group basis to protect the system from self-interested behaviour [8].

2.3 Information Flow in Social Networks

The dissemination of information is a ubiquitous process in human social networks and computer systems. It plays a fundamental role in numerous settings including the penetration of technological innovations [29], the word-of-mouth effect [22], the propagation of news and opinion [18], and distributed problem-solving [21]. The principle underlying how information diffuses from person to person mimics the spread of an epidemic or a rumour [11], expanding to a substantial proportion of the population in a short number of steps according to the ‘small-world’ phenomenon [30].

In our recent work we analysed many aspects of information flow in social networks. Firstly we have been concerned with the mathematical modelling of social selection whereby individuals tend to form relationships with others who are already similar to them. We explored two distinct dynamics that yield such phenomena: social proximity [13, 14] and social utility [15] in the context of opportunistic communication networks. Moreover we have investigated the process of social influence that leads people to adopt behaviours exhibited by those they interact with. We have analysed the impact of the logical structure of social networks on the spread of rumours [12]. We have also investigated the dynamics of the word-of-mouth effect through networks of people [1].

The two forces of social influence and selection are common place in a wide range of social settings: when deciding to adopt an activity individuals imitate the people they are currently interacting with; and they simultaneously form new interactions as a result of their existing activities. These two aspects of social networks are crucial to our understanding of how people form and update their opinions by interacting with a small sample of the population, namely their neighbours.

3 MICRO-SOCIAL SYSTEMS

Based on this background, in this section, we give an informal definition and brief discussion of the concept of micro-social systems.

The motivation for micro-social systems stems from the commonplace observation that networked computing devices are the driving force of modern industry, entertainment and commerce, providing powerful, and inter-connected, infrastructures for business organizations (through eCommerce, holonic manufacturing and agile enter-

prises), computer-mediated communication, and mobile ad hoc networks (i.e. MANETs, vehicular ad hoc networks (VANETs), etc.).

Research in networked computing also needs to take into consideration the three following features:

- Local information, partial knowledge and inconsistent union: What each network node ‘sees’ is the result of actions by (possibly millions) of actors, some of which are not known, and even those actions which are known, the actors motive may be unknown. Moreover, what a node ‘thinks’ it ‘sees’ may not be consistent with the ‘opinion’ of other nodes.
- Decentralised control: there is no single central authority that is controlling or coordinating the actions of others. The emphasis is on local decision-making based on locally available information and the perception of locally witnessed events.
- Social organization: from this it follows that in the absence of perfect knowledge there is no perfect form of government, therefore the next best thing is a government prepared to modify its policies according to the needs, requirements and expressed opinions of its ‘population’.

In other words, social organization is both the requirement for and consequence of any networked computing which impacts on personal, legal or commercial relationships between real-world entities (people or organizations). The social sciences, together with legal and organizational theory, provide rich models of social organization, from which we can ‘cherry pick’ appropriate concepts which can be formalised in corresponding logical and computational models. These computational models are, to some extent, a scaled-down or simplified version of the original social system or theory – hence micro-social system.

Therefore we define (informally) a micro-social system as a distributed computer system or network where the interactions, relationships and dependencies between components is a microcosm of aspects of a human society. We will use the term ‘agent’ to denote one of these components, both to indicate membership of such a social system (e.g. as a node in a network) and in the computational sense of ‘software agent’, i.e. as an encapsulated, embedded, autonomous process responsible for its own state and decisions.

The aspects of human society in which we are interested includes communication protocols, organizational rules and hierarchies, network structures, inter-personal relationships, and other processes of self-determination and self-organization. In particular, though, we pick out the following three primary set of rules which underpin the social intelligence required to realise a micro-social system which can be applied to the issues in ad hoc networks detailed earlier (i.e. dynamism, conflicts, sub-ideality, security and continuity):

- Rules of Social Order. Micro-social systems (MSS) consist of agents whose actions have a conventional significance (according to the social rules of an institution); actions are therefore norm-governed. This requires characterising the permissions, obligations and (institutional) powers of each agent to determine which actions are valid, meaningful, legal, and so on.
- Rules of Social Choice. MSS consist of heterogeneous, self-interested agents that can have conflicting preferences in decision-making situations; these preferences can be aggregated by taking votes over potential outcomes. In practice, an election is held, and the winning candidate is declared to be the agreed choice.
- Rules of Social Exchange. MSS, being both open and local, will require agents to gain knowledge over time by exchanging information with each other. Each agent must therefore be capable of

reliable opinion formation, based on the opinions gathered from the contacts in their own social networks. Processes of belief revision, belief merging, judgement aggregation and truth tracking are therefore important.

The development of a micro-social system is illustrated in Figure 1. We start from first principles, and analyse specific aspects of existing networked societies, i.e. human social systems, namely rules of social order, social exchange and social choice (i.e. rules, interactions and group decision-making). From this we derive the formal framework for specification of micro-social systems, which we apply to design, implement and run a computational system. Moreover, we aim visualise the operation of such systems in the same, human-understandable terms (There remains the possibility of using micro-social systems to simulate human societies, but we do not address this topic further in the current work.)




Figure 1. Developing micro-social systems

One of the key aspects of this development cycle are the mechanisms for run-time (self-)adaptation of the micro-social system. In particular, we seek to go beyond emergent behaviour seen, for example in swarm intelligence for emergent systems, i.e. the non-introspective application of hard-wired local computations, with respect to physical rules and/or the environment, which achieve unintended or unknown global outcomes, to the social intelligence for micro-social systems, which involves the introspective application of soft-wired local computations, with respect to physical rules, the environment and conventional rules, in order to achieve intended and coordinated global outcomes.

This is illustrated in Figure 2. Note that the formation of an organization involves hierarchies of groups, each with their own roles (chair, head, etc.), conventional rules (institutional regulations), constituency, remit (as defined by some ontology), and so on.




Figure 2. Formation and Adaptation of Organizations

Many other issues are raised by this form of self-organization, but in particular, we need to understand: firstly, how an arbitrary collection of agents can self-organise into an organisational structure; how an arbitrary collection of agents can self-organise its social network (which is structurally distinct from the organizational structure); and thirdly, what is the interplay between the explicit formal organization and the implicit social network. We illustrate this issue in the

next section, with a discussion of a scenario, developing an adaptive security model for mobile ad hoc networks.

4 ADAPTIVE NETWORK SECURITY

In this section, we discuss an application scenario for micro-social systems, in the context of developing an adaptive network security policy for mobile ad hoc networks (MANETs).

4.1 Problem Description

Users access networked information and services from a variety of mobile devices over a variety of ad hoc networks. The devices have a wide range of characteristics, from processor speed to battery power; while the networks are open and volatile, implying that the network nodes cannot necessarily be trusted, either to be present or even to operate correctly. For applications to operate in a non-trusted computer network, there are many different security mechanisms, which differ in terms of cost, algorithmic complexity, types of license, and so on. Similarly there are many proposed trust/reputation frameworks, which also vary in the computational complexity of computing the trust decision. In a resource-constrained, i.e. power challenged, computing environment, It is essential to choose the most appropriate security mechanisms to match the features of content sensitivity, service delivery, device typology, and network topology.

This choice is too fast, frequent and complex for users: they will simply ignore security altogether, or worse default to the strongest level of security, to the detriment of resource utilisation, network throughput and quality of service. Instead, the choice should be delegated to the devices and applications themselves; but clearly there is a trade off involved, between algorithmic complexity (and so power cost), and the level of security and/or the (potential) accuracy of the trust decision. We propose to implement this choice via an Adaptive Network Security (ANS) scheme using the framework of micro-social systems. In the next two sections, we look at two models, one for rules of social exchange, and another for rules of social order (specifically for adaptation).

4.2 Social Exchange

In specifying a model of social exchange, we start from social science theories, such as [16, 17], which we formalise using graph theory, and represent ideas of influence, opinion and affinity to define a formal model. The properties of this model, e.g. in polarisation, fragmentation, and consensus, etc., are investigated in [28].

4.2.1 Social Network

We define the agents as the nodes and *confidence relations* between agents as the edges of a social network. We define a social network at timepoint t as a directional, weighted graph, $G = (N, R)_t$, consisting of a non-empty set of agents N of cardinality n , and an incidence relation $R \subseteq N \times N \times [0, 1]$.

Each agent is characterised by:

- A task:** purpose for which the agent was designed. A task defines what the agent should achieve while being part of the system.
- A strategy:** action(s) which the agent follows in order to fulfil its task.
- A mind-set:** hides the *true preference* of the agent about the issue under discussion.

An opinion: communicates the *expressed preference* of the agents about the issue under discussion.

A level of self-confidence: confidence that an agent has in its own opinion. This value might change as new opinions are received.

A level of confidence on others' opinions: weights the relation that an agents has in each of its acquaintances.

4.2.2 Modeling Confidence

From R , for any given timepoint t , we can derive a family of confidence functions $W = \{w_1, w_2, \dots, w_n\}$ where each function w_i is of the form $w_i : N \times T \rightarrow [0, 1]$. Accordingly, each confidence function is expressed as $w_i(j, t)$. For simplicity, we use the reduced form $w_{i,j}(t)$ to express the confidence (function) that agent i assigns to agent j at time t .

Thus, each confidence function $w_{i,j}(t)$ assigns a real value between 0 and 1 to the confidence relation between the ordered pair $\langle i, j \rangle$, indicating how much confidence i has in j at a specific time point $t \in T = \{0, 1, \dots\}$. When $j = i$ the confidence function $w_{i,i}(t)$ yields a measure of *self-confidence*.

Following the well-established tradition in the field of opinion formation (see e.g. Friedkin [17] and Hegselmann [20]), we consider that for each agent the sum of the confidence in its acquaintances is always 1 (including itself), $\sum_{j=1}^n w_{i,j}(t) = 1$. This ensures that the measure of confidence is not based on absolute inter-agent judgements but on relative intra-agent ones, and each agent's relative confidence in the others increases and decreases only with respect to its own value judgements.

4.2.3 Modeling Opinion Formation Dynamics

Starting from the simplest assumption that agents are able to communicate very simple, concrete pieces of information with one another, and that one agent might receive this information from different sources, we are concerned with the problem of aggregating information from many sources into one.

The type of information that we deal with in this study is that of subjective nature. We conceptualise information as "information-as-knowledge" as defined by Buckland in [6]: "A key characteristic of "information-as-knowledge" is that it is intangible: one cannot touch it or measure it in any direct way. Knowledge, belief and opinion are personal, subjective, and conceptual". We refer to "information-as-knowledge" as opinions.

Each agent holds information in the form of opinions, $o_i : T \rightarrow [0, 100]$. We adopt a continuous opinion approach, in line with [10, 17, 20], and consider an agent i 's opinion at time t , $o_i(t)$, as a real-valued representation of an opinions (rather than a specific measure).

We assume that each agent holds an initial opinion (i.e. $o_i(1) \neq \perp$) on every issue about to be discussed. However this opinion can change with time as agents are influenced by opinions exchanged with other agents. The influence that one agent's opinion exerts on another's is given by how much confidence, relative to other agents, the latter has in the former.

Correspondingly, the opinion formation dynamics consists of simultaneous opinion exchanges between pairs of agents and a subsequent individual opinion revision. The main objective of an agent is to collect opinions from other agents in order to revise (i.e. consolidate or modify) its own opinions. A secondary objective is to share its own opinions and influence other agents towards them.

The opinion formation dynamics occurs at discrete time points and on a per issue basis. At each time point each agent exchanges opin-

ions 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) such that:

$$o_i(t + 1) = \sum_{j=1}^n w_{i,j}(t) o_j(t) \quad (1)$$

4.2.4 Modeling Affinity Between Agents

We specify the matching between an agent i 's mind-set and another agent's opinion by defining an *affinity function* $a_i : N \times T \rightarrow [0, 1]$. This function evaluates the linear similarity between an opinion and a given constant μ which is a representative reference value of an agent's *mind-set* for a given issue. Correspondingly, we express the affinity function as $a_i(j, t)$. Again for simplicity, we denote this affinity (function) between i 's mind-set and j 's opinion as $a_{i,j}(t)$ and we define it as:

$$a_{i,j}(t) = 1 - \frac{|o_j(t) - \mu_i|}{\max(\mu_i, 100 - \mu_i)} \quad (2)$$

μ_i is agent i 's mind-set, which is constant per issue per agent, although no two agents need to have the same mind-set on any issue. Thus in each time step, the *affinity* between agents can be different for each ordered pair of agents corresponding to the fitness between opinions and *mind-sets*.

4.2.5 Modelling Social Influence

One of the main characteristics of our model is that we assume that agents rely differently on other agents. Thus agents can have more confidence in some agents than others and this can change with time. Agents start the opinion formation protocol with an initial value of confidence in each of the agents within their social network. As the first opinion exchange happens and the opinion of each agent changes according to equation 1, the confidence changes accordingly.

In our model, an agent i increases its confidence in another agent j based on how well j 's opinion coincides with i 's *mind-set*. Assuming a positive evaluation for those opinions matching agent i 's *mind-set* 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 *mind-set*.

Therefore, confidence changes in time differently for each agent, based on the *affinity* between agents. Agents increase the confidence in those agents whose opinions fit their *mind-set*. Thus the confidence in other agents is redistributed according to the following equation:

$$w_{i,j}(t + 1) = \frac{w_{i,j}(t) + w_{i,j}(t) a_{i,j}(t)}{\sum_{k=1}^n (w_{i,k}(t) + w_{i,k}(t) a_{i,k}(t))} \quad (3)$$

Equation 3 contains the principle that an increase in confidence in those agents with an initially given high confidence value is higher than in those with an initially given low confidence value.

4.3 Social Order

We can now use the rules of social exchange in conjunction with rules of social order as the basis for adaptation, providing mechanisms with which the specification of a micro-social system may be modified during its execution (cf. [3]).

System participants may require to change at run-time their institutional powers, permissions, obligations, rights, etc. We envisage a system architecture composed of several protocol levels: the protocol at the bottom level (0-level protocol) corresponds to the procedure for conducting the business of the system – for example, in a MANET the 0-level protocol could be a security protocol.

At any point in time during the execution of the 0-level protocol the participants may start a 'meta' protocol in order to modify the rules of the 0-level protocol. A meta-protocol can be a procedure for voting, that is, the system participants vote for or against a proposed modification in the rules of the 0-level protocol. Moreover, the participants of the meta protocol may initiate a meta-meta protocol to modify the rules of the meta protocol, or they may initiate a meta-meta-meta protocol to modify the rules of the meta-meta protocol, and so on. In general, assuming an infrastructure of maximum k protocol levels, the protocol participants of a n -level protocol ($0 \leq n \leq k$) may start a $n + m$ -level protocol ($1 \leq m \leq k$) in order to modify the rules of the $n + m - 1$ -level protocol (see Figure 3).




Figure 3. Protocol Stack for Rules of Social Order

The key problem is to judge precisely when to start the 'meta' protocol. Recall, that in a power-challenged environment, there is scant value in each agent re-assessing the 'fitness' of the network given the current security policies and protocols vs. the power-cost to implement them. Rather, we use the rules of social exchange to filter the opinion "it is necessary to change" or "it is not necessary to change" through the network. For example, the security policy and protocols at the base level (level 0) may be traded off against the available power and/or the trustworthiness of the nodes in the network: at low power but with fully trusted nodes, a low-level security policy may be appropriate; at 'high' power with a risk that some nodes may have been compromised by some network attack, a higher-level security policy may be required.

Depending on the population profile and the model of social exchange, as defined in the previous section, at some point, if an opinion leader starts to advocate the opinion that "it is necessary to change", this will be disseminated throughout the network (in a computable time), will modify the opinions of other agents, the pressure of which (i.e. the current state being X and enough agents 'think' (express the opinion) that X is unsatisfactory), and eventually induce a self-confident agent to initiate the meta-protocol.

5 SUMMARY AND CONCLUSION

We have introduced the concept of *micro-social systems*, scaled down agent societies based on a computational microcosm of human soci-

eties. Micro-social systems are an attempt to bring together principles of social intelligence – including rules of social order and rules of social choice – with principles of social networks – including rules of social exchange, and other factors such as proximity, influence, and utility.

We propose to use micro-social systems as a formal, computational framework on which to design and specify agent societies for various types of ad hoc network, whether computer-mediated communication, mobile networks, or virtual organizations. We considered a scenario involving an adaptive network security scheme, in which rules of social order and rules of social exchange were interleaved to prompt proposals for run-time self adaptation. This self-organization could be applied to develop an adaptive network security scheme, for example by changing the security protocols and procedures to suit the available power and trustworthiness of the nodes. A more substantive model than that presented here would accommodate a wider range of security policy conflicts, e.g. following the comprehensive classification of [19].

It is evident then that this work is at a formative stage and there are many details yet to be worked out in the interleaving of rules of social order and rules of social exchange, not to mention rules of social choice. When this is formulated, we can start applying the framework to deal with other issues, in particular conflicts. Conflicts are inevitable in open systems; however experience in (human) legal and commercial systems shows that those systems which facilitate internal resolution of conflicts (i.e. without recourse to time-consuming and expensive litigation) emerge with stronger social and inter-personal relationships. Therefore one immediate avenue of further research is the interleaving of rules of social order for argumentation and conflict resolution protocols [5, 27] with rules of social exchange used in juries to reach consensus and decisions. Rules of social choice are also of particular significance here, to prevent manipulation of the outcome.

In terms of experimentation, we propose to use the PreSAGE multi-agent simulation platform [24], which provides substantial support for investigating the effect of agent interactions, network properties and organizational rules on individual agent behaviour and long-term collective global performance.

ACKNOWLEDGEMENTS

We would like to thank Alexander Artikis (Demokritos, Athens) for his contribution to the research.

REFERENCES

- [1] H. Amini, M. Draief, and M. Lelarge, 'Marketing in a random network', in *Springer LNCS, Network Control and Optimization*, p. To appear, (2008).
- [2] A. Artikis, L. Kamara, J. Pitt, and M. Sergot, 'A protocol for resource sharing in norm-governed ad hoc networks', in *Second International Workshop on Declarative Agent Languages and Technologies (DALT'04)*, eds., J. Leite, A. Omicini, P. Torroni, and P. Yolum, volume 3476, pp. 221–238, (2005).
- [3] A. Artikis, D. Kaponis, and J. Pitt, 'Dynamic specifications of norm-governed systems', in *Multi-Agent Systems: Semantics and Dynamics of Organisational Models*, ed., V. Dignum. IGI Global, (2009).
- [4] A. Artikis and J. Pitt, 'Specifying open agent systems: A survey', in *Engineering Societies in the Agents World: Proceedings ESAW08*, eds., G. Picard, L. Vercouter, and A. Artikis, volume LNAI, p. to appear. Springer Verlag, (2009).
- [5] A. Artikis, M. Sergot, and J. Pitt, 'Specifying norm-governed computational societies', *ACM Transactions on Computational Intelligence*, **10**(1), to appear, (2009).
- [6] M.K. Buckland, 'Information as thing', *Journal of the American Society for Information Science*, **42**(5), 351–360, (1991).
- [7] H. Carr and J. Pitt, 'Adaptation of voting rules in agent societies', in *Proceedings AAMAS Workshop on Organised Adaptation in Multi-Agent Systems*, eds., G. Vourros, A. Artikis, K. Stathis, and J. Pitt, volume LNAI5368, p. to appear. Springer Verlag, (2008).
- [8] H. Carr, J. Pitt, and A. Artikis, 'Peer pressure as a driver of adaptation in agent societies', in *Proceedings AAMAS Workshop on Organised Adaptation in Multi-Agent Systems*, eds., A. Artikis, G. Picard, and L. Vercouter, p. to appear. Springer Verlag, (2008).
- [9] Y. Chevalere, U. Endriss, J. Lang, and N. Maudet, 'A short introduction to computational social choice', *SOFSEM: Theory and Practice of Computer Science*, 51–69, (2007).
- [10] Guillaume Deffuant, David Neau, Frédéric Amblard, and Gérard Weisbuch, 'Mixing beliefs among interacting agents', *Advances in Complex Systems*, **3**(1–4), 87–98, (2000).
- [11] P. Dodds and D. Watts, 'Universal behavior in a generalized model of contagion', *Physical Review Letters*, **92**, 218701, (2004).
- [12] M. Draief, 'Epidemic processes on complex networks', *Physica A: Statistical Mechanics and its Applications*, **363**(1), 120–131, (2006).
- [13] M. Draief and A. Ganesh, 'Efficient routing in poisson small-world networks', *Journal of Applied Probability*, **43**(3), 678–686, (2006).
- [14] M. Draief and A. Ganesh, 'Spread of epidemics and rumours with mobile agents', arXiv:0810.3128, http://eprintweb.org/S/authors/All/dr/Draief/1.* (under revision).
- [15] M. Draief, A. Ganesh, and L. Massoulié, 'Exponential random graphs as models of overlay networks', arXiv:0810.3173, <http://eprintweb.org/S/authors/All/dr/Draief/2> (under revision).
- [16] J. French, 'A formal theory of social power', *Psychological Review*, **63**, 181–194, (1956).
- [17] Noah E. Friedkin and Eugene C. Johnsen, 'Social influence and opinions', *Journal of Mathematical Sociology*, **15**, 193–205, (1990).
- [18] D. Gruhl, D. Liben-Nowell, R.V. Guha, and A. Tomkins, 'Information diffusion through blogspace', in *Proceedings of the 13th International World Wide Web Conference*, eds., S.I. Feldman, M. Uretsky, M. Najork, and C.E. Wills, pp. 491–501. ACM Press, (2004).
- [19] H. Hamed and E. Al-Shaer, 'Taxonomy of conflicts in network security policies', *IEEE Communications Magazine*, **44**(3), 134–141, (2006).
- [20] Rainer Hegselmann and Ulrich Krause, 'Opinion dynamics and bounded confidence: models, analysis and simulation', *Journal of Artificial Societies and Social Simulation*, **5**(3), (2002).
- [21] M. Kearns, S. Surit, and N. Monfort, 'An experimental study of the coloring problem on human subject networks', *Science*, **313**, 824–827, (2006).
- [22] J. Leskovec, L. Adamic, and B. Huberman, 'The dynamics of viral marketing', in *Proceedings of the 7th ACM Conference on Electronic Commerce*, eds., J. Feigenbaum, J.C.I. Chuang, and D.M. Pennock, pp. 228–237. ACM Press, (2006).
- [23] B. Neville and J. Pitt, 'A computational framework for social agents in agent mediated e-commerce', in *AAMAS Trust Workshop*, ed., R. Falcone, pp. 83–91, (2004).
- [24] B. Neville and J. Pitt, 'Presage: A programming environment for the simulation of agent societies', in *Proc. AAMAS Workshop on Programming Multi-agent Systems (ProMAS)'08*. Springer Verlag, (to appear).
- [25] E. Ostrom, *Governing The Commons: The Evolution of Institutions for Collective Action*, Cambridge University Press, 1990.
- [26] J. Pitt, L. Kamara, M. Sergot, and A. Artikis, 'Voting in multi-agent systems', *Computer Journal*, **49**(2), 156–170, (2006).
- [27] J. Pitt, D. Ramirez-Cano, L. Kamara, and B. Neville, 'Alternative dispute resolution in virtual organizations', in *Proceedings ESAW 2007*, eds., A. Artikis, G. O'Hare, K. Stathis, and G. A. Vourros, volume 4995 of *Lecture Notes in Computer Science*, pp. 72–89, (2007).
- [28] D. Ramirez-Cano and J. Pitt, 'Follow the leader: Profiling agents in an opinion formation model of dynamic confidence and individual mindsets', in *Proceedings of the 2006 IEEE/WIC/ACM International Conference on Intelligent Agent Technology*, pp. 660–667, (2006).
- [29] E. Rogers, *Diffusion of Innovations*, Free Press (4th Edition), 1995.
- [30] J. Travers and S. Milgram, 'An experimental study of the small world problem', *Sociometry*, **32**, 425–443, (1969).
- [31] A. Vasalou, J. Pitt, and G. Piolle, 'From theory to practice: Forgiveness as a mechanism to repair conflicts in cmc', in *Proceedings iTrust 2006*, eds., K. Stølen, W. Winsborough, F. Martinelli, and F. Massacci, volume LNCS3986, pp. 397–411. Springer, (2006).