Report of Dagstuhl Seminar 15482

Social Concepts in Self-organising Systems

Edited by

Ada Diaconescu[1], Stephen Marsh[2], Jeremy Pitt[3], Wolfgang Reif[4], and Jan-Philipp Steghöfer[5]

1 **Telecom Paris Tech, FR,** ada.diaconescu@telecom-paristech.fr
2 **UOIT – Oshawa, CA,** stephen.marsh@uoit.ca
3 **Imperial College London, GB,** j.pitt@imperial.ac.uk
4 **Universität Augsburg, DE,** reif@informatik.uni-augsburg.de
5 **Chalmers UT – Göteborg, SE,** jan-philipp.steghofer@cse.gu.se

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 15482 "Social Concepts in Self-organising Systems". The seminar brought together researchers from computer sciences (in particular from the fields of multi-agent systems and self-organisation) and from social sciences to discuss the impact of the use of social concepts in technical systems as well as the interaction between technical and social systems. In an engaging and interactive setting, the problem was illuminated from a technical as well as a philosophical and legal point of view. The talks, discussions, and working groups identified a growing body of work in the field, a number of interesting and promising research avenues, as well as a set of open issues for future investigation.

Seminar November 22–27, 2015 – http://www.dagstuhl.de/15482
1998 ACM Subject Classification H.1.2 User/Machine Systems, I.2.11 Distributed Artificial Intelligence, K.4 Computers and Society
Keywords and phrases computational justice, multi-agent sytems, norms, organic computing, reputation, self-organisation, social capital, socio-technical systems, sociologically-inspired computing, trust
Digital Object Identifier 10.4230/DagRep.5.11.127

1 Executive Summary

Jan-Philipp Steghöfer
Ada Diaconescu
Stephen Marsh
Jeremy Pitt
Wolfgang Reif

License Creative Commons BY 3.0 Unported license
© Jan-Philipp Steghöfer, Ada Diaconescu, Stephen Marsh, Jeremy Pitt, and Wolfgang Reif

There are two exciting trends in computing that motivated this seminar. On the one hand, large-scale self-organising systems gain traction in real-world settings, e.g., in the autonomous control of the power grid or in personal transportation scenarios. On the other hand, our lives are more and more pervaded by socio-technical systems that rely on the interaction of existing, complex social systems and technical systems that in many ways mirror and form the social relationships of their users. The seminar brought together researchers from a variety of domains to discuss the technical, legal, and social issues these trends incur. One focus was how social concepts can be formalised and implemented to make technical

Except where otherwise noted, content of this report is licensed
under a Creative Commons BY 3.0 Unported license

self-organising systems more robust and efficient. The other focus was how technology shapes the social system and vice versa.

Use of Social Concepts in Self-Organising and Socio-Technical Systems

The seminar's first focus is motivated by the requirements of large-scale self-organising systems. The more such systems have to take their environment into account, the more open, and the more heterogeneous they are, the more important social concepts become [1]. If a population of agents is no longer developed, deployed, maintained, and controlled by a single company or institution, the goals of the agents no longer concur—especially, the individual sub-goals no longer necessarily imply the overall system goal. In such cases, social constructs can help encourage cooperation between the agents. The presence of norms as an explicit expression of acceptable behaviour [3], of a trust management system to encourage reciprocity [4], or of a form of computational justice to settle disputes within the system [5] are measures that have been discussed in the scientific community for these ends. In this way, the systems form a legal reality that establishes certain rules and regulations within the system. This legal reality must be in accordance with the legal system under whose jurisdiction these systems work.

The second focus is how technical systems interact with and influence existing social systems. With the increasing dependence of society on computation and on complex artificial systems, their influence on human-computer interaction, and on inter-human interaction becomes a topic of concern (see, e.g., [2]). One aspect of this is the novel challenge of managing an online identity, made necessary by the representations of human users in technical systems that are, necessarily, an abstraction of the real user. Another aspect is the increasing reliance of human users on these technical aids and the potential of negative effects on the users accompanied with this. Such effects can range from infringement of privacy, to withholding of relevant information, and even to targeted manipulation.

Results of the Seminar

The seminar was highly interactive, with a lot of time dedicated to plenum discussions. Talks were used as impulses to stimulate discussion and working groups focused on particular aspects that the participants deemed particularly important.

A number of talks addressed the implications of using social concepts in technical systems from different angles and featured insights into existing technical solutions, e.g., for computational justice, trust, and ethical behaviour, as well as the observable effects of these solutions. Likewise, incentives and how social constructs influence them was a recurring theme. The discussions following these talks addressed important issues such as the relationship between system and user values, goals and the rules designed to achieve these goals, possible attacks on socio-technical systems, quantifiable incentives, and self-determination in technical systems.

A different set of talks was aimed at understanding the way social systems and technology interact, e.g., how the social organisation of the human users is represented in the technical system and becomes evident in the interactions in that system. An overview of the interplay of legal and technical systems was also provided, with important insights into the connection between technical feasibility and legal admissibility. This set of talks encouraged discussion geared towards governance, power, and the representation of values in technical systems, as well as on how to represent existing social systems in technical systems and how both the social system and its representation evolve over time.

Based on the discussions and the input from the talks, three working groups were formed that focused on discussing different aspects of the use of social concepts in self-organising systems. Their main aims and contributions are as follows:

Understanding: A first step was to consider what the notion of "social" means in the context of the technical systems regarded in the seminar. Based on a brief literature survey, the working group determined that social means that an organisation exists, that the welfare of the individuals and the organisation is regarded, and that the relations between individuals and between individuals and the organisation are a concern. A second step was to stipulate that formalising social values leads to the individuals behaving in a way that recognises their social obligations and responsibilities. Finally, the notion of "socially-sensitive design" was introduced to denote that both the design process and the system itself must be socially sensitive.

Engineering: The main concerns were how to make different social concepts usable in technical systems, how to select the fitting social construct for a specific problem, how to measure its effectiveness, and how to combine several social concepts. The working group suggested a pattern language to express selection criteria, implementation approaches, and consequences, as well as a set of metrics that make it possible to evaluate the impact of the social concept in the technical system.

Dynamics: The discussion developed towards how the social and technical components of socio-technical systems interact with each other and how the resulting dynamic aspects influence these systems. A total of six challenges were identified, the most important of which pertains to how the interaction between different social concepts that provides "checks and balances" in social systems can be transferred to technical systems. Further problems that were discussed are conflict resolution and power distribution, as well as the influence of technical systems on society and where the responsibility for this influence lies.

Future Work

The seminar participants agreed that the topic is timely and relevant and that there are a number of open issues that need to be addressed in the future. Possible venues for future elaboration of these issues are the SASO^ST workshops[1], held annually at the IEEE Conference on Self-Adaptive and Self-Organising Systems, as well as a number of other projects currently in discussion. In particular, the organisers are discussing ways to provide an overview of the state of the art of the field as well as a research roadmap and opportunities to specifically discuss the impact self-organising and socio-technical systems will have on society.

References

1. L. Vercouter and G. Muller. L.i.a.r.: Achieving social control in open and decentralized multiagent systems. *Applied Artificial Intelligence*, 24(8):723–768, 2010.
2. J. Botev. Anonymity, immediacy and electoral delegation in socio-technical networked computer systems. In M. Horbach, editor, *Informatik 2013, 43. Jahrestagung der Gesellschaft für Informatik e.V. (GI), Informatik angepasst an Mensch, Organisation und Umwelt, 16.-20. September 2013, Koblenz*, volume 220 of *LNI*, page 1164. GI, 2013.
3. T. Balke, C. da Costa Pereira, F. Dignum, E. Lorini, A. Rotolo, W. Vasconcelos, and S. Villata. Norms in MAS: Definitions and related concepts. In G. Andrighetto, G. Governatori,

[1] http://sasost.isse.de/

P. Noriega, and L. van der Torre, editors, *Normative Multi-Agent Systems*, volume 4 of *Dagstuhl Follow-Ups*, pages 1–31. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2013.
4 L. Mui, M. Mohtashemi, and A. Halberstadt. A computational model of trust and reputation. In *Proc. of the 35th Hawaii International Conference on System Sciences (HICSS'02)*, pages 188–196. IEEE Computer Society Press, 2002.
5 J. Pitt, D. Busquets, and R. Riveret. Formal models of social processes: The pursuit of computational justice in self-organising multi-agent systems. In *7th IEEE International Conference on Self-Adaptive and Self-Organizing Systems, SASO 2013, Philadelphia, PA, USA, September 9-13, 2013*, pages 269–270. IEEE, 2013.

2 Table of Contents

3 Overview of Talks

3.1 When is an Interaction . . . well just an interaction

Kirstie Bellman (Topcy House Consulting, US)

License Creative Commons BY 3.0 Unported license
© Kirstie Bellman

The Space Station has two and a half million moving parts, which does not include the non-moving parts, the complex launch vehicle, the ground stations, the special test equipment built to support its development and continual refinement. The coordination of components, the establishment of interfaces and communications among components and the integration of required behavior from the behavior of massively different components is a major part of the decade long development of space systems. The relationships among components can be dynamic, even semi-autonomous, but they are not social. The purpose of this talk is to discuss the need for teasing apart our concepts of integration, interfaces, collective behavior, and social processes, with an emphasis on what we gain from these different levels of integration and socially-aware collective behaviors.

*with nod to Freud in title

3.2 Society of autonomous agents & Ethics

Olivier Boissier (Ecole des Mines – St. Etienne, FR)

License Creative Commons BY 3.0 Unported license
© Olivier Boissier

The increasing use of multi-agent technologies in various areas raises the necessity of designing agents and societies of agents that produce ethical behaviors in context. Considering the development of socio-technical systems where humans delegate part of their decisions to agents and where user-centric approaches are required, we investigate the dimensions to be considered in order to define systems able to exhibit ethical behaviors at runtime. We claim that besides individual reasoning mechanisms and representations at the micro/agent level, representations and mechanisms have also to be investigated at the macro/society level.

3.3 Anchoring Institutions: Intermediate Institutions and the Meta-Rule of Law

Pompeu Casanovas (Autonomus University of Barcelona, ES)

License Creative Commons BY 3.0 Unported license
© Pompeu Casanovas

How to bridge institutions, rules, norms, apps and people to set up specific ecosystems that turn “legal”? How to regulate “legally” the information flow on the Web in order to empower people (individuals and communities) and make the balance between liberty and security? How to make “legally” effective artificial devices – electronic institutions, Rights Expression Languages (REL. . .) – on the Web of Data? This presentation tries to provide some ways to answer these questions. Law has changed its traditional meaning on the Web of Data.

It requires the construction of intermediate institutions – e.g., Semantic Web Regulatory Models (SWRM) – to anchor the reuse of ontologies, datasets, and general knowledge into specific contextual legal settings. At present, there are coordination and cohesion problems between the regulatory instruments – law, policies, standards, and ethical principles – and the scenarios set forth within the Web of Data. I furnish three different examples: (i) the regulation of CAPER (a European platform to fight organised crime); (ii) lessons learned in the making of the Catalan White Book on Mediation; (iii) Licensed Linked Data. All three can be faced as cases of Relational Law.

3.4 (Socially-inspired Technology-reinforced) (Values and Pathologies)

Ada Diaconescu (Telecom Paris Tech, FR)

License Creative Commons BY 3.0 Unported license
© Ada Diaconescu

This presentation aims to explore the mutual relation between social and technical systems, which are both instances of self-adapting self-organising systems. The presentation raises several questions related to the opportunities and risks that may occur in the context of socially-inspired technical systems, which are executing and interacting within a social environment. Opportunities include the social values that such systems can help achieve or reinforce, even in large-scale highly-dynamic societies (different from "traditional" ones). Challenges include the social pathologies that technical systems might import along with the social concepts and processes, and which they may reinforce and magnify. Additional challenges may occur due to the lack of sufficient understanding when modelling social concepts; and/or to the context discrepancies between the social context where they evolved initially and the technical environment where they are imported. Finally, technical systems may disturb or disrupt existing social processes by altering social structures (e.g., via new communication links) and their dynamics (e.g., via reduced time scales for communication and adaptation).

3.5 Toward incentives for self-organization in a decentralized data-sharing system

Babak Esfandiari (Carleton University – Ottawa, CA)

License Creative Commons BY 3.0 Unported license
© Babak Esfandiari
Joint work of Esfandiari, Babak; Davoust, Alan
Main reference A. Davoust, A. Craig, B. Esfandiari, V. Kazmierski, "P2Pedia: a peer-to-peer wiki for decentralized collaboration," Concurrency and Computation: Practice and Experience, 27(11):2778–2795, 2015.
URL http://dx.doi.org/10.1002/cpe.3420

To quote Tim Berners-Lee, "we need to re-decentralize the web". This is because the web is in the hands of central authorities who can exercise control over the contributions of users, through censorship and biases in search result rankings. They can also simply disappear, taking away the entirety of user contributions. We want to evaluate an alternative where authority is decentralized among users. For such a system to work, there needs to be incentives for users to contribute and to self-organize. Producers of documents derive their payoff from the appearance of their contribution in consumer search results. Consumers derive their payoff through the relevance of the ranking of their search results. We propose ranking metrics that, while helping consumers maximize their payoff, also incentivize their participation in the system by mirroring documents and managing their network neighborhood. The diversity of these metrics also helps make the system more resilient to trust attacks (sybil, social exploitation). We sketch a simulation that helps us validate these claims.

3.6 All Human Values are System Values

Stephen Marsh (UOIT – Oshawa, CA)

License Creative Commons BY 3.0 Unported license
© Stephen Marsh
Joint work of Marsh, Stephen; Dwyer, Natasha; Basu, Anirban

Sociotechnical systems are those that embody values, of their designers and creators and users – as well as the other agents in the system. To an extent, the way in which those values are represented is less important than the recognition that they exist, can be represented, and an be used for, for example, comparison (the "Jensen Question"). In this talk I examine the importance of values such as trust, peace and forgiveness, and provide Ten Commandments, originally applied to trust management, that I hope can help in the design and actualisation of all sociotechnical systems.

The commandments were first presented in:
Stephen Marsh, Anirban Basu, Natasha Dwyer: Rendering unto Cæsar the Things That Are Cæsar's: Complex Trust Models and Human Understanding. IFIPTM 2012:191–200.

Later they were increased (to 10) in an exploration in:
Stephen Marsh, Natasha Dwyer, Anirban Basu, Tim Storer, Karen Renaud, Khalil El-Khatib, Babak Esfandiari, Sylvie Noël, and Mehmet Vefa Bicakci: Foreground Trust as a Security Paradigm: Turning Users into Strong Links, Information Security in Diverse Computing Environments, 8 pages, 2014, IGI Global.

3.7 Governance, Sustainability and Justice

Jeremy Pitt (Imperial College London, GB)

License Creative Commons BY 3.0 Unported license
© Jeremy Pitt

Many open comuting systems – for example grid computing, sensor or vehicular networks, or virtual organisations – face a similar problem: how to collectivise and distribute resources in the absence of a centralsied controller. On approach is to define a set of conventional, mutable and mutually-agree rules – ie a self-organising rule-oriented systems in which the rules as explicit first-class entities, typically characterised by an institution. In this talk weconsider the use of self-organisig rule-oriented systems based on formal models of Ostrom's institutional design principles and Rescher's theory of distributive justice as a basis for inclusive and sustainable alloation of common-pool resources. We identify the social concepts of governance (decision-making procedures underpinning operational, collective and constitutional choice rules), sustainability (as the underlying goal of the procedures) and justness (correctness in the outcome of those procedures).

3.8 Interactional Justice

Jeremy Pitt (Imperial College London, GB)

License Creative Commons BY 3.0 Unported license
© Jeremy Pitt

We present interactional justice as a way to increase the 'correctness' or 'appropriateness' of outcomes of algorithmic decision- making and deliberative processes in self-organising multi-agent systems, paving the way for their effective operation in both decentralised networks and user-centred socio-technical systems. We investigate how the social concept of interactional justice can be formalised in computational logic, to understand better principles of fairness in resource allocation, inclusivity in self-determination, and fitness of procedural rules. As a result, self-organising systems can be designed and deployed for a wide range of applications, from ad hoc networks to community energy systems, wherein qualitative (social) values are primary system requirements.

3.9 Understanding Social Organizations: A Network Perspective

Ingo Scholtes (ETH Zürich, CH)

License Creative Commons BY 3.0 Unported license
© Ingo Scholtes
Main reference I. Scholtes, "Understanding complex systems: When Big Data meets network science," it – Information Technology, 57(4):252–256, 2015.
URL http://dx.doi.org/10.1515/itit-2015-0012

The convergence of social and technical systems provides us with a wealth of data on the structure and dynamics of social organizations. It is tempting to utilize these data to better understand how social organizations evolve, how their structure is related to their "performance", and how the position of individuals in the emerging social fabric affects their motivation. Taking a network perspective on such questions, in this talk I will introduce recent research results obtained in the context of empirical software engineering. They demonstrate the potential of computational methods in the study of social phenomena and mechanisms. At the same time, I will highlight fallacies arising in the application of the complex networks perspective to complex socio-technical systems.

References

1 Ingo Scholtes, Pavlin Mavrodiev and Frank Schweitzer: *From Aristotle to Ringelmann: A large-scale analysis of productivity and coordination in Open Source Software projects*, to appear in Empirical Software Engineering, Springer, November 2015
2 Marcelo S. Zanetti, Ingo Scholtes, Claudio Juan Tessone and Frank Schweitzer: *Categorizing Bugs with Social Networks: A Case Study on Four Open Source Software Communities*, In Proceedings of the 35th International Conference on Software Engineering (ICSE '13), pp. 1032-1041, San Francisco, CA, USA, May 2013
3 Marcelo Serrano Zanetti, Ingo Scholtes, Claudio Juan Tessone and Frank Schweitzer: *The Rise and Fall of a Central Contributor: Centralization and Performance in the Gentoo Community*, In Proceedings of the 6th International Workshop on Cooperative and Human Aspects in Software Engineering (CHASE 2013) held at ICSE 2013, San Francisco, CA, USA, May 2013

4 Emre Sarigöl, René Pfitzner, Ingo Scholtes, Antonios Garas and Frank Schweitzer: *Predicting Scientific Success Based on Coauthorship Networks*, In EPJ Data Science, Springer, September 25 2014
5 Ingo Scholtes: *Understanding Complex Systems: When Big Data meets Network Science*, In it – Information Technology, Methods and Applications of Informatics and Information Technology, August 1 2015

3.10 A few experiences on using trust for social control

Laurent Vercouter (INSA – Saint-Étienne-du-Rouvray, FR)

License Creative Commons BY 3.0 Unported license
© Laurent Vercouter

The social concept of trust has been widely used to implement social control for distributed systems. It is particularly suited to develop local trust assessment algorithms to evaluate the trustworthiness of an agent's acquaintances and to use it in collective decision processes for system adaptation or reconfiguration. This talk studies trust model variations defined for different use cases: peer-to-peer, social or sensor networks. We emphasize the specific characteristics of each of these cases that justifies an adaptation of the concepts or the models to obtain a relevant mapping into a (socio-)technical system. More specifically, the integration of human users in a network or the lack of an identity management system have a major impact on the way social control is implemented.

4 Working groups

4.1 Working group on Engineering Social Concepts

Olivier Boissier (Ecole des Mines – St. Etienne, FR), Gerrit Anders (Universität Augsburg, DE), Babak Esfandiari (Carleton University – Ottawa, CA), Gauthier Picard (Ecole des Mines – St. Etienne, FR), Wolfgang Reif (Universität Augsburg, DE), and Laurent Vercouter (INSA – Saint-Étienne-du-Rouvray, FR)

License Creative Commons BY 3.0 Unported license
© Olivier Boissier, Gerrit Anders, Babak Esfandiari, Gauthier Picard, Wolfgang Reif, and Laurent Vercouter

The working group about *engineering social concepts for self-organizing systems* has been interested in the study of the contributions of social concepts for technical and socio-technical systems and in the way they can be mapped into these systems. The use of social concepts for self-organizing systems, and more specifically socio-technical systems raises several questions and challenges from the engineering point of view.

The need for self-organizing systems arises when the system to be designed has to live in a changing and hard to predict environment, with no centralized control, and where the participating entities are potentially autonomous (this includes humans), heterogeneous and/or unreliable. Human societies share some of these properties, and the social concepts they use present the potential for being helpful metaphors and abstractions, just as ethology and biomimetism has already provided inspiration for design paradigms. Inclusiveness is another use case. When the human is in the loop, there is a need for the system to explain itself and interact using the same concepts as the ones humans use. Privacy is an area where this need seems relevant.

There are several examples of technical or socio-technical systems in which social concepts are useful. Future power management systems are one of them. They are characterized by an increasing number of small power producers, such as biogas power plants owned by farmers, and consumers, such as individual households. To ensure the system's stability and efficiency, these entities have to play an active role in the system. This imposes several challenges: Usually, power markets define a threshold for the minimum production or consumption (hereinafter referred to as prosumption) and participants have to guarantee that they can provide a specific prosumption over a certain amount of time. Furthermore, due to the shift to renewable energy sources and due to the consumers' stochastic behavior, future prosumption is subject to uncertainty. These uncertainties not only have to be taken into account to identify reliable trading partners but also to decide about the volume of a contract, for instance. There are several other application examples in which social concepts are useful in order to deal with problems such as uncertainty, the presence of human users, or the need to have decentralized and adaptive algorithms. We can cite wireless sensor networks, desktop grid computing systems, or applications in the context of smart production and smart cities, among others.

There are many challenges when attempting to translate "soft" concepts into computing systems. The concept first needs to be observed, understood and described. Then it needs to be formalized to the extent that its definition is not ambiguous. At this stage, there is the risk of leaving out crucial aspects due to oversimplification. Next, using directly the formalization as a computable model may prove intractable, and so further simplifications, translation and application to the problem domain may even reduce the usefulness the concept further. Finally, verification and validation of a technical solution may be challenging due to the lack of proper metrics and benchmarks.

In the process of engineering these systems, we encounter recurring problems and solutions based on these social metaphors. We propose to use a pattern based approach like design patterns in classical software engineering (this approach has already been used in the multi-agent domain, see [1]). These patterns can then hopefully be incorporated into a "social toolbox" which would provide us with a framework or an API supporting built-in social primitives. The self-organizing systems built with such a toolbox should then be evaluated using benchmarks that would measure the same metrics that where guiding the choice of the social pattern. These patterns are described by four fields corresponding to (i) a problem; (ii) a solution; (iii) consequences; (iv) examples.

Besides these patterns, it is necessary to have metrics and evaluation tools to emphasize the contribution of socially-based solutions. Several kinds of metrics are needed in order to evaluate the impacts (advantages/drawbacks) of introducing such concepts and mechanisms in the system. Metrics could be organized as follows depending on their functional or non functional aspects.

- Metrics on Functional Requirements
  - Application metrics, i.e. metrics related to the application domain targeted by the system
  - Social mechanism metrics, i.e. metrics related to the evaluation of the social mechanisms built from the considered social metaphors
- Metrics on Non Functional Requirements
  - Technical metrics, e.g. scalability, tractability, resilience, time to repair (mean time to failure)

- Social metrics, i.e. user acceptance, e.g. willingness to give up autonomy for a social benefit as for instance in the consumer energy area, trustworthiness of consumers such as failure to comply leading to paying a price. Dedicated benchmarks are also required so that we can make experimental comparisons of different solutions to a given problem, and thus be able to determine the relative efficiency of different social concepts or different models inspired by a same concept. Some benchmarks have already been produced in the past, such as the ART Testbed [2] for trust models assessments. These experiences have shown that the essential aspects for an efficient benchmark are:
  - to find a common scenario simulating typical problems addressed, using the social concepts we target;
  - to define specific metrics.

At last, we can emphasize a few research challenges while applying social concepts to self-organizing systems. A first one is to be able to define a computer model from an observed social phenomena going through steps of pre-formal and formal specifications as explained in the beginning of this section. Another challenge is to define a social concept pattern language to specify solutions (patterns, relations between patterns, etc.). Then specific patterns may be defined by considering a set of social concepts and a set of self-organizing system problems to find out where and how a social concept may bring an interesting solution to a given problem.

References

1 Kendall, Elizabeth A., et al. Patterns of intelligent and mobile agents. Proceedings of the second international conference on Autonomous agents. ACM, 1998.

2 Karen Fullam, Tomas B. Klos, Guillaume Muller, Jordi Sabater, Andreas Schlosser, Zvi Topol, K. Suzanne Barber, Jeffrey S. Rosenschein, Laurent Vercouter, Marco Voss, A specification of the Agent Reputation and Trust (ART) testbed: experimentation and competition for trust in agent societies. AAMAS 2005: 512–518

4.2 Dynamics in and of Social Concepts in Self-organising Systems

Sebastian Götz (TU Dresden, DE), Nelly Bencomo (Aston University – Birmingham, GB), Pompeu Casanovas (Autonomus University of Barcelona, ES), Ada Diaconescu (Telecom Paris Tech, FR), Jan Kantert (Leibniz Universität Hannover, DE), Christian Müller-Schloer (Leibniz Universität Hannover, DE), Jan-Philipp Steghöfer (Chalmers UT – Göteborg, SE), and Leon van der Torre (University of Luxembourg, LU)

License Creative Commons BY 3.0 Unported license
© Sebastian Götz, Nelly Bencomo, Pompeu Casanovas, Ada Diaconescu, Jan Kantert, Christian Müller-Schloer, Jan-Philipp Steghöfer, and Leon van der Torre

The discussion on social concepts in self-organising systems can be approached based on two central questions: (a) how to transfer concepts and approaches from social sciences to technical systems and (b) how are social, technical and socio-technical systems interrelated. While the first question was addressed in the working groups on understanding and engineering, this working group focused on the second question and tailored the discussion towards the dynamic aspects of the relation between social systems and technical systems. These dynamic aspects are evident whenever processes in one of the related systems have an impact on the other. For a discussion on terminology and the engineering process to be used, we refer the interested reader to the summaries of the two respective working groups.

The main research objective of this working group was to identify a list of challenges, which are to be addressed in socio-technical systems. These challenges are intended to fuel future research and can hopefully act as the basis for a research roadmap. The following six challenges emerged from the discussions.

The first challenge denotes the need to identify various approaches for managing the mutual influence between technical and social systems.

As technological developments become ever faster, the frequency of changes to technical systems also increases. Moreover, as technology itself becomes faster and more self-adaptive, the pace and frequency of change perceivable when interacting with modern technology are even greater. This, in turn, can raise problems for socio-technical systems, since the technical parts allow for very fast responses, while the social parts respond much slower. For example, if TripAdvisor introduces a new feature such as a social network for the users, how does that change the social system constituted by its users? While the technical change is instantaneous from the perspective of the users, the actual growth of the social network and the emergence of social interactions through the technical feature can take a considerable time, if it emerges at all. In addition, such a technical service can impact the social structure of the users, their behaviour, and their interaction with the system. Additionally, since users become more reachable anywhere, at any time, via communication technology, they may experience increasing peer-pressure for more reactivity and availability. The quality of the communication may also be impacted.

A promising first step is to think about how a technical system can be integrated within a change process of the social system (e.g., as defined in a company). For this, a change process should be specified that describes how the technical system can adapt in unison with the social system. An interesting phenomenon to discuss in the context of this challenge is the *micro-macro feedback loop*. This loop denotes that the behaviour of individuals (micro-level) has an impact on a higher-level behaviour, or property, such as the overall socio-technical system, or its environment (macro-level); in turn, this higher-level phenomenon (macro-level) also has an effect on the individuals (micro-level). To approach the aforementioned challenge, this loop of mutual effects has to be considered carefully.

The second challenge is about the dangers of importing social concepts into socio-technical systems without importing the checks and balances that are part of the complex social system. Cultural evolution has produced a set of interlinked systems that provide some level of resilience to the overall collective (or system of systems) with respect to faulty or malicious behaviours from individual participants. These interlinked systems also provide means of excluding participants temporarily and of reintegrating them later on if their behaviour improves – e.g., by applying the social concept of forgiveness. The collective *together* produces a relatively stable, fair and just social system. However, in technical systems, we use isolated concepts in order to make quick decisions, which can have dire consequences for the agents who have little possibility of recourse.

For instance, the specification and enforcement of norms within a system requires that a special-purpose agent role is defined in addition, to verify the manner in which such norm-related processes are being carried-out, and hence to make sure that the associated powers are not being abused. In many legal systems, such control is, e.g., provided by the "checks and balances" achieved through the separation of powers. Likewise, when introducing a trust and reputation system, forgiveness must become part of the system to avoid the isolation of participants. Therefore, introducing new concepts from social systems may be seen like opening a Pandora's box, since it becomes necessary to import more and more social concepts to ensure that the ones the system designers originally had in mind actually work as intended.

Since self-adaptive and self-organising systems can, theoretically, change arbitrarily, and since certain behaviours can contradict the system's objectives, or inconvenience and disrupt other systems, or humans, system behaviour generally needs to be constrained, or regulated. In normative systems for instance, norms are defined to regulate, constrain and/or attach meaning to the actions of agents within the system. Here, rules and laws are made-up and agreed-upon by convention; and special-purpose mechanisms set in place to monitor and enforce them, as well as to correct non-compliant behaviour [3]. For instance, in social systems, human agents are designated to certain roles, which gives them institutional power. They can then perform acts which create conventional, mutually-agreed facts (e.g., declaring that a couple is "married"). An important question to consider here is who is allowed (or empowered) to create these facts [4]. And, in extension "Qui costudiet ipsos custodes"? Currently, in technical systems, mechanisms for ensuring checks-and-balances and conflict-resolution are most often external to the system – e.g. relying on the "traditional" legal system for settling issues in technical systems. This approach features several limitations, including the different time-scales at which technical systems and "traditional" social systems operate, and evolve – e.g. the traditional legal system lags behind modern socio-technical systems, and hence no longer addresses resulting necessities. This disregards one of Ostrom's principles [2] which requires the availability of conflict-resolution mechanisms that are fast and efficient. Additional problems may be caused if the normative system, which is external to a targeted socio-technical system, exercises too much power and hence limits the self-regulation autonomy of the socio-technical system. This also contravenes with one of Ostrom's design principles – related to the right to self-organise – and may in turn lead to the failure to manage common resources in a sustainable way (e.g., sharing electric power in a smart grid).

Since norms, laws and rules are conventional facts, rather than physical ones, and, therefore, can be broken, important challenges will be raised concerning enforcement mechanisms and their limitations. The next two challenges aim to highlight some examples of these.

The third challenge is therefore about conflict resolution, which is one of Ostrom's principles that aims to avoid the negative runaway dynamics that lead to the tragedy of the commons [1]. In other words, there is a need for inherent complexity in the regulations that police conflict resolution. For instance, the regulations on delays and reimbursements of passengers of airlines are comparatively simple, but leave a lot of room for interpretation and conflict (e.g., when should one consider that "the airplane is late"?). Nonetheless, who is doing the conflict resolution? Is this entity fair, or trustworthy? Are the resolvers empowered to act? Do the agents accept the resolution of the conflicts?

The fourth challenge identified in this working group covers the problem of imbalanced power distribution in socio-technical systems. For example, when there is a resource allocation problem and one agent feels unfairly left out of the allocation, by the time the conflict is resolved, this agent might have "starved", especially if the conflict resolution mechanism is external and lengthy. This denotes an imbalanced distribution of power, since the party withholding the resources has power over the starved individual.

Apart from the challenges above that relate directly into how to incorporate social concepts into technical systems and the consequences this might have, a fifth one deals with accountability for these consequences, both positive and negative.

The fifth challenge is concerned with the question of who takes responsibility for the impact of socio-technical systems – e.g. upon society, the environment and so on. One option would be to only hold social scientists responsible, since this would fall within their area of expertise; while computer scientists would merely be required to think about how to actually build the technical systems. In this case, lawyers would need to regulate the systems, based on advisory input from social scientists on the human-related aspects, and from computer scientists on the technical risks and limitations.

On the other hand, it could also be that computer scientists should share part of the responsibility and aim to figure-out how their systems might impact social systems, on various time scales. The precautionary principle should not be ignored here, even if it is in fact routinely ignored in practice. In this scenario, discussions and agreement among computer scientists, social scientists and lawyers would be required.

The final challenge the working group addressed is, again, a more technical one and one that is directly related to the capabilities technical systems that interact with social systems can provide.

The sixth and final challenge is about the possibility of continuous optimisation. Due to the constant interaction between the technical and the social systems it becomes possible to continuously measure the social system, adapt the technical system to the social system very quickly and, thus, to make it more suitable for its users. This optimisation could also occur for the goals of the owners or for another specific group that has different interests from the original user group.

References

1 Garrett Hardin. *The Tragedy of the Commons*. Science 162 (3859), pp. 1243–1248, 1968.
2 Elinor Ostrom. *Governing the commons: The evolution of institutions for collective action*. Cambridge University Press, 1990.
3 Andrew Jones and Marek Sergot. *On the characterisation of law and computer systems: The normative systems perspective*, Deontic logic in computer science: normative system specification, p. 275–307, 1993
4 Andrew Jones and Marek Sergot. *A formal characterisation of institutionalised power*. Journal of the IGPL, 4(3), pp. 429–445, 1996.

4.3 Socially-Sensitive Systems Design – Working Group on "Understanding Social Concepts in Self-Organising Systems"

Peter R. Lewis (Aston University – Birmingham, GB), Kirstie Bellman (Topcy House Consulting, US), Jean Botev (University of Luxembourg, LU), Hanno Hildmann (NEC Laboratories Europe – Heidelberg, DE), Stephen Marsh (UOIT – Oshawa, CA), Jeremy Pitt (Imperial College London, GB), Ingo Scholtes (ETH Zürich, CH), and Sven Tomforde (Universität Augsburg, DE)

License Creative Commons BY 3.0 Unported license
© Peter R. Lewis, Kirstie Bellman, Jean Botev, Hanno Hildmann, Stephen Marsh, Jeremy Pitt, Ingo Scholtes, and Sven Tomforde

4.3.1 Introduction

This Working Group addressed the question of how to understand and formalise the role of social concepts in self-organising systems. Further, to achieve systems which are aware of, or in other ways sensitive to social concepts, we will be required to be able to design social concepts in to technical systems, in a principled way. In considering systems which explicitly contain social concepts (or an awareness of them), we proposed the notion of *socially-sensitive* *systems*. We then further proposed steps towards the creation of a conceptual framework for the design of socially-sensitive systems.

In a society, individuals from all walks of life voluntarily organise themselves in groups to gain benefits which improve their quality of life. There are however issues associated with free riding, scale, and time. We argued that by appealing to socially-sensitive systems design, we can support individuals by obtaining a sufficient social position to be resilient to these issues of scale and time.

Some of the new concepts discussed in socially-sensitive systems design will include how going beyond the notion that collective behaviour is only driven by immediate goals to now have concepts and mechanisms for more enduring, value-driven group behaviour and for incorporating social mechanisms supporting human social values.

One of the key benefits of socially-sensitive systems design is therefore that it leads to systems which engage in better positioning through the increase of social potential. This implies continuous redesign of the social aspects of the system, in order to ensure the resilience of the system as a whole, in a way that generalises to unknown situations. We discussed several early scenarios for how agent based system could include concepts and mechanisms for both joint goal-behaviour and for supporting the group, leading to fewer social pathologies and more robust systems.

The proposed conceptual framework for socially-sensitive systems design has three tenets: social organisation, social values and social relations. Social organisation relates to issues of network structure and roles. Social values relate to *states that matter* to individuals, and can through social mechanisms come to matter to the group as a whole. For social relations, we build on Sztompka's [9] sociological hierarchy of social relations, which differentiate a spectrum from social behaviours through social actions, to social interactions and social relations.

In this Working Group, we proposed that socially-sensitive systems design has the potential to lead to systems in which individuals behave in a way that recognises their social obligations and responsibilities. Further, we argued that this ensures the endurance of the social aspects of the system, as well as the benefits they bring.

4.3.2 Why Design Socially-Sensitive Systems

Socially-sensitive systems and design will bring benefits in two quite different types of system. In purely technical systems, a key benefit will be derived from the notion of *better positioning through increasing social potential*. Uncertainty often exists around what individuals or a group as a whole may face, or will be required to do, especially as complexity increases. Nevertheless, entities within a system will still need to achieve certain goals, often quickly. Often in a complex system, this will require interaction with, or even perhaps cooperation of, other entities within the broader system. As one example of this, resolving a resource contention issue with a degree of immediacy may require other entities to give up claims to that resource quickly. Further, these types of challenges are typically not one-offs: entities within a system may know, or learn, that they are going to have to iterate. The other entities involved in the interaction, in our example perhaps the one who is asked to give up the resource, will be encountered again. In future uncertain scenarios, perhaps the pattern will be repeated, or the roles reversed. Given the uncertainties associated with these complex systems, operating in unfolding situations, there will be a need to be able to account for and generalise to future unknown situations such as this.

In these cases, groups of individuals use social organisation, values and relations to move themselves *as a group*, to a better position to be able to deal with these factors. In doing so, the group builds what we call *social potential*[2] Without social concepts explicitly embedded in a system and its decision-making processes, there will be a lack of primitives with which to reason about the system's *social state*, and increased social potential cannot be explicitly targeted.

One example of this in the human sphere is a football team who, while they may not be under immediate attack from the opposition, nor know what form such an attack may take, apply organisation (in the form of roles and positions), values (e.g. solidarity, respect, fairness and loyalty [7]) and relations [6] to put themselves in a better position to deal with such an attack, when it does come. On a longer timescale, such social positioning has been found not only to increase the potential for success within the specific game or even group itself, but this even generalises to future life events, such as professional success [6].

However, the rise of socio-technical systems and their weaving into the fabric of human social interaction means that such values will be ever more important. Indeed, we even venture as far as to argue that such an approach is not only desirable, it is essential to preserve the intrinsic richness and value of human social interaction, as technical systems are increasingly interwoven into our everyday lives. To do otherwise, would be to degrade the quality of the human experience in a socio-technical world.

4.3.2.1 Issues With Existing Approaches

We anticipate that it will be useful to more formally understand the impact of different forms of organisation, and the role of various forms of social relation, in socio-technical systems. However, we identify a major issue with existing approaches being that of how to incorporate and manage the value we, as humans, associate with social aspects. In many traditional approaches, for example in much of multi-agent systems, an approach is generally taken whereby direct (often numerical) comparability of alternatives may be assumed, based on a universal commodification of such values.

We argue that this approach is, in general, insufficient to capture human social values such as obligation, empathy, peace and justice. Nevertheless these are the things that often truly matter the most to humans and human society. What is needed is a proper set of theories for how to actualise these things in computational systems. Such theories will be essential to realise the vision of socially-sensitive systems and design.

Our aim is to produce computational formalisations of these social values, enabling them to be explicitly represented within the technical side of socio-technical systems. In doing so, we provide some of the tools required to ensure that such systems also uphold these social values.

4.3.3 A Conceptual Framework for Socially-Sensitive Systems Design

When we talk of systems or a design process being *socially sensitive*, we mean specifically that a system is both *socially aware* and that it is *socially active*. Social awareness implies that the system can observe social aspects of its environment and interactions within in, and conceptualise these, in order to reason about social aspects. Further, *socially active* implies that a system does not simply observe and think, but based on its conceptualisations, acts in a way that is congruent with them, and its own social principles.

[2] This is not the same as social potential identified in [8], but is more closely related to social capital, in the spirit of how Fukuyama [5] sees it.

We proposed a framework for reasoning about socially-sensitive systems and their design. This is built upon three tenets:

1. The group's social organisation, including the network structure, individuals' roles and perhaps rank within it.
2. The group's social values, specifically preferences associated with states of the group and its individuals.
3. The group's social relations, both in terms of type and structure of relations.

4.3.3.1 Social Organisation

Organisation is perhaps the most familiar feature of the social nature of technical systems. It is concerned with network structure, roles of individuals and sub-groups within the system, and other features such as rank.

There is now a substantial literature on the organisation and self-organisation of (socio-)technical systems, and these aspects will underlie many of the other social concepts which a system may explicitly possess. Indeed, the organisation might be thought of as the platform, or set of constraints, upon which social relations play out, and social values are observed and propagated.

4.3.3.2 Social Values

Unpacking the notion of social values, we can relate the behaviour of individuals within a collective to obligations and responsibilities. Agents no longer care only for the goals of the group, but also for the members of the group themselves, in terms of their values for how they care for each other. In general, we can consider social values to be descriptions of *states that matter* to individuals, and can through social mechanisms come to matter to the group as a whole. More concretely, we might consider:

- States of a group that matter,
- States of an individual that matter,
- States of a relationship that matter.

States that matter to individuals can matter to groups too. This may be realised through a variety of mechanisms, such as collective decision-making and aggregation.

As discussed above, some types of value properties will be economic, insofar that they relate to things that can be readily quantified, or at least compared, without losing their primary essence. But other value properties, those that we call *welfare values*, describe qualities of the system that are less readily quantifiable or comparable. These express preferences concerning the ways in which thing are done, in accordance with what matters to individuals and groups. One potential way for formalising this notion is through the use of meta-goals, or constraints over meta-goals.

4.3.3.3 Social Relations

Weber [10] claimed that an action is 'social' if the acting individual takes account of the behaviour of others and is thereby oriented in its course. Further, Sztompka proposed [9] a hierarchy of social interactions and relations, shown in Table 1. The hierarchy makes clear how many social concepts already familiar to computer scientists relate to each other. For example, *social behaviour* is commonly analysed in ant-based systems (e.g. [2, 3]), where actions (e.g. leaving pheromone) are done for the benefit of other ants, and in doing so for the benefit of the colony as a whole. Similarly, *regular interactions* describe the kinds of

Table 1 Sztompka's Sociological Hierarchy of Social Relations.

Type	Requires
Behaviour	Physical movement
Action	Meaning
Social Behaviour	Directed towards others
Social Action	Await response
Social Contact	Unique / rare interaction
Social Interaction	Interactions
Repeated Interaction	Accidental, not planned, but repeated interaction
Regular Interaction	Regularity
Regulated Interaction	Interactions described by law, custom or tradition
Social Relation	A scheme of social interactions

interactions occurring in repeated games such as the iterated prisoners' dilemma [1]. Here, knowledge of future interactions with the same opponent is crucial to determining future behaviour.

However, Sztompka's hierarchy demonstrates that there are many forms of social interaction, which require varying forms of knowledge (concerning both oneself, other individuals, and the environment) as well as cognitive capabilities. In developing this tenet of our framework, our intention is to map Sztompka's framework to computational systems, and extend as necessary.

4.3.3.4 Relationship to Value-Sensitive Design

As is clear from the presented framework, in addition to a sensitivity to organisation and relations, socially-sensitive systems must be sensitive to values. Of course, this implies the practice of approaches such as value-sensitive design [4]. But socially-sensitive systems go beyond this, not just being designed by designers (or co-designed) in accordance with social values. The systems themselves will also be sensitive to such values. Thus, a form of socially-sensitive meta-design is needed. Indeed, due to additional complexities present in socio-technical self-organising systems, such as unexpected dynamics and continuous reorganisation, we may even require socially-sensitive self-design, as the system plays an active role in its own design, on a continuous basis, in accordance with social values that it itself promotes.

4.3.4 Conclusions

In summary, we propose that for socio-technical systems to possess, be aware of, and act in accordance with social concepts, these social concepts will need to be formalised and made explicit. Further, we argue that they will need to be designed in. We describe such systems and their design process as *socially-sensitive systems design*. A key benefit of socially-sensitive systems will be better positioning, through increased social potential.

While there is much work needed to formalise and realise the notion of socially-sensitive systems design, it is clear at this stage that any list of requirements for socially-sensitive systems will include at the very least:

- groups,
- awareness of others,
- directed behaviour, and
- cognition.

In this Working Group, we proposed a framework for the socially-sensitive systems design, based on three core tenets of social organisation, social values and social relations.

Ultimately, we anticipate the benefit to technical systems to include increased robustness, increased empathy with humans, a reduction of pathologies of digital communities. Further, there is also the potential for insights gained in building and using socially-sensitive systems to impact on our understanding of human society itself. As a result of this understanding feeding back into social science, we expect that we can better support human beings in society at large.

In continuing the work from this Working Group, we plan to further develop the conceptual framework which we have sketched here. This will first include mapping and possibly extending Sztompka's social relations hierarchy for computational systems. Second, we will relate existing work on individual computational values (e.g. trust) to the framework. Third, we will look to formalise other social concepts, which are not yet, or only partially explicitly present in technical systems. These include those things that often really matter to humans, such as obligation, justice and peace.

References

1 R. Axelrod. The evolution of strategies in the iterated prisoner's dilemma. In L. Davis, editor, *Genetic algorithms and simulated annealing*, pages 32–41. Pittman, London, 1987.

2 Marco Dorigo. *Optimization, Learning and Natural Algorithms*. PhD thesis, Politecnico di Milano, Milan, Italy, 1992.

3 Lukas Esterle, Peter R. Lewis, Xin Yao, and Bernhard Rinner. Socio-economic vision graph generation and handover in distributed smart camera networks. *ACM Transactions on Sensor Networks*, 10(2), 2014.

4 Batya Friedman and Peter H. Kahn Value sensitive design: Theory and methods. Technical Report 02-12-01, Deptartment Of Computer Science and Engineering, University of Washington, 2002.

5 Francis Fukuyama. *Trust: The Social Virtues and the Creation of Prosperity*. Free Press, 1996.

6 A Huggins and S Randell. The contribution of sports to gender equality and women's empowerment. In *Proceedings of the International Conference on Gender equity on Sports for Global Change*, 2007.

7 Esther Rutten, Geert-Jan Stams, Gert Biesta, Carlo Schuengel, Evelien Dirks, and Jan Hoeksma. The contribution of organized youth sport to antisocial and prosocial behavior in adolescent athletes. *Journal of Youth and Adolescence*, 36(3):255–264, 2007.

8 John H. Reif and Hongyan Wang. Social potential fields: A distributed behavioral control for autonomous robots. *Robotics and Autonomous Systems*, 27(3):171–194, 1999.

9 Piotr Sztompka. *Socjologia*. Znak, 2002.

10 Max Weber. *The Nature of Social Action*. Cambridge University Press, 1978.

5 Open problems

5.1 Your Cheating Cat!

Stephen Marsh (UOIT – Oshawa, CA)

License Creative Commons BY 3.0 Unported license
© Stephen Marsh
Main reference The talk is based on a corpus of the work I've done previously.

Technical Systems are becoming more opaque, for various reasons, some valid an some less so. However, for we (human's and otherwise) who are using or are affected by them, this is an issue. I'm beginning to explore ways to mitigate this through the various strengths of the technologies themselves. This wee ideas talk will explore the problem, think about these strengths, and try to spark discussion around our options and possible work, including in ethics, monitoring and morality.

5.2 The concept of self-reorganization

Gauthier Picard (Ecole des Mines – St. Etienne, FR)

License Creative Commons BY 3.0 Unported license
© Gauthier Picard
Joint work of Picard, G., Persson, C., Boissier, O., and Ramparany, F.
Main reference G. Picard, C. Persson, O. Boissier, F. Ramparany, "Multi-agent self-organization and reorganization to adapt M2M infrastructures," in Proc. of the 9th IEEE Int'l Conf. on Self-Adaptive and Self-Organizing Systems (SASO'15), pp. 91–100, IEEE CS, 2015.
URL http://dx.doi.org/10.1109/SASO.2015.17

Designing and monitoring complex systems raise major challenges, due to the multitude of heterogeneous components, which have their own internal dynamics, while in interaction with a highly dynamic and uncertain environments. In such cases, centralised management is not realistic and predefined system behaviors lead to obsolescence. So, how to equip systems with bounded autonomous adaptation capabilities to handle these complexities and dynamics? (By bounded, we mean we want to keep control on the system, by constraining its behaviour.) From a multi-agent engineering perspective, we translate this question into "How to set up multi-agent organisation adaptation process enabling the emergence of consistent and desired behaviours to develop adaptive systems?"

We propose to join forces coming from self-organisation approaches (e.g. swarms) which exhibits important adaptiveness capabilities, and reorganisation (e.g. organisational multiagent systems) making explicit the structure, the functionalities and the constraints on the organisation. This join approach is coined "self-reorganisation". Based on multiagent programming approach JaCaMo, we share our experience on the implementation of self-reorganising systems in smart city, ambient intelligence and trust management application fields.

5.3 Why would anybody want to change?

Jan-Philipp Steghöfer (Chalmers UT – Göteborg, SE)

License © Creative Commons BY 3.0 Unported license
© Jan-Philipp Steghöfer

Change management theories recognise that organisational inertia and individual resistance are detriments to change. In this talk, I briefly describe these notions and point out the similarities to self-organising systems and their potential relevance for socio-technical systems in particular. Finally, I pose the question of whether we must consider the ability and willingness to change in agent-based systems and what the impact on self-organisation will be.

References

1 B. Burnes, "Kurt Lewin and the Planned Approach to Change: A Re-appraisal," Journal of Management studies, vol. 41, no. 6, pp. 977–1002, 2004.
2 K. E. Weick and R. E. Quinn, "Organizational Change and Development," Annual Review of Psychology, vol. 50, no. 1, pp. 361–386, 1999.
3 J. P. Kotter, "Leading Change: Why Transformation Efforts Fail," Harvard Business Review, vol. 73, no. 2, pp. 59–67, 1995.
4 Katherine J. Klein and Joann S. Sorra, "The Challenge of Innovation Implementation," The Academy of Management Review, 21(4):1055–1080, 1996.
5 Kurt Lewin, "Frontiers in Group Dynamics: Concept, Method and Reality in Social Sciences: Social Equilibria and Social Change", Human Relation, vol. 1, no. 36, 1947

Participants

- Gerrit Anders
Universität Augsburg, DE
- Kirstie Bellman
Topcy House Consulting, US
- Nelly Bencomo
Aston Univ. – Birmingham, GB
- Olivier Boissier
Ecole des Mines –
St. Etienne, FR
- Jean Botev
University of Luxembourg, LU
- Pompeu Casanovas
Autonomus University of
Barcelona, ES
- Ada Diaconescu
Telecom Paris Tech, FR
- Babak Esfandiari
Carleton Univ. – Ottawa, CA
- Sebastian Götz
TU Dresden, DE
- Hanno Hildmann
NEC Laboratories Europe –
Heidelberg, DE
- Jan Kantert
Leibniz Univ. Hannover, DE
- Peter R. Lewis
Aston Univ. – Birmingham, GB
- Stephen Marsh
UOIT – Oshawa, CA
- Christian Müller-Schloer
Leibniz Univ. Hannover, DE
- Gauthier Picard
Ecole des Mines –
St. Etienne, FR
- Jeremy Pitt
Imperial College London, GB
- Wolfgang Reif
Universität Augsburg, DE
- Ingo Scholtes
ETH Zürich, CH
- Jan-Philipp Steghöfer
Chalmers UT – Göteborg, SE
- Sven Tomforde
Universität Augsburg, DE
- Leon van der Torre
University of Luxembourg, LU
- Laurent Vercouter
INSA –
Saint-Étienne-du-Rouvray, FR