INVITED TALK
Logics of Interaction, Coalitions and Social Choice (extended abstract)

Thomas Ågotnes[1 2] and **Wiebe van der Hoek**[3] and **Michael Wooldridge**[4]

Abstract. While different forms of social interaction have been extensively studied in several fields, the development of formal logics makes precise knowledge representation and mechanical reasoning about situations involving social interaction possible. In particular, such logics make it possible to formally specify and verify software implementing social mechanisms. In my talk I will give an overview of some of our recent work on logics for social interaction, in particular applied to reasoning about social choice mechanisms such as voting and preference aggregation as well as reasoning about coalition formation and coalitional stability. We use benchmark examples from game theory and social choice theory to illustrate the expressiveness of the logics.

1 Introduction

Logics for reasoning about social interaction take into account the facts that individual agents are *autonomous*, *self interested*, and act *strategically* in order to obtain their goals. Such logics often try to give an account of what happens when agents choose to *cooperate* or otherwise act together [7, 17]. Of course, concepts such as strategic interaction, preferences, and cooperation have been extensively studied in fields such as game theory and social choice theory. However, the study of their *logical* and *computational* properties is relatively new.

Logics for social interaction are useful for, e.g., knowledge representation and reasoning in multi-agent systems [21] and for the formal specification and automated verification of *computational mechanisms* [19, 20, 13]. Given a specification of a mechanism in the form of a logical formula, we can use logical tools to *verify* (model checking) and/or *synthesise* (constructive proof of satisfiability) mechanisms [18, 23].

In this note we briefly discuss some new logics for social interaction. Many such logics have been suggested, and they differ, first and foremost, in their *expressiveness*. In order for a logic to be useful it must be sufficiently expressive, and we take concepts, properties and results from fields such as game theory and social choice theory as benchmark tests of expressiveness.

In this note we shall be particularly concerned with two aspects of interaction. The first is *social choice mechanisms*; a very general class of economic mechanisms [9] concerned with selecting some particular outcome from a range of alternatives on behalf of a collection of agents, such as *voting*, *preference aggregation* and *judgment aggregation*. The second is *coalition formation* in general and *coalitional stability* in particular. While it is clear that these are two distinct notions and problems, they are closely connected at least on one level of abstraction: they are both intimately related to the concept of *coalitional power* – what coalitions of agents can make come about. In the remainder of this note we very briefly and informally introduce some new logics for reasoning about these concepts, illustrated with example formulae (see the references for details). As a starting point we first mention Marc Pauly's seminal logic of coalitional power: Coalition Logic.

2 Coalition Logic

Two popular logics for social interaction are Alur, Henzinger and Kupferman's *Alternating-time Temporal Logic (*ATL*)* [7] and Pauly's *Coalition Logic (*CL*)* [17]. These logics let us express properties about the abilities, or powers, of single agents and of groups of agents acting together. The main syntactic construct of Coalition Logic is of the form[5]

$$\langle C\rangle\varphi,$$

where C is a set of agents, or a *coalition*, and φ is a formula. The intended meaning of $\langle C\rangle\varphi$ is that the agents C can choose to make φ come about, by performing some joint action. If $\langle C\rangle\varphi$ is true, then φ will not necessarily come about (C might choose to not make φ come about), but if $\langle C\rangle\varphi$ is true then C can guarantee that, no matter what the other agents do, φ will come about. ATL adds temporal operators such as "sometime in the future" to the language as well (CL can be seen as the next-time fragment of ATL).

Formally, formulae of Coalition Logic are interpreted on state-based structures where the agents play a strategic game, in the sense of non-cooperative game theory[6], in each state.

Pauly also observed [16] that the Coalition Logic construct can be used to express properties of social choice mechanisms. Consider the following example of a simple social choice mechanism [16]:

> *Two individuals, A and B, must choose between two outcomes, p and q. We want a mechanism that will allow them to choose which will satisfy the following requirements: We want an outcome to be possible – that is, we want the two agents to choose, collectively, either p or q. We do not want them to be able to*

[1] In my talk I will present joint work with Wiebe van der Hoek and Michael Wooldridge. In this extended abstract we give a brief outline.
[2] Bergen University College, Norway, email: tag@hib.no
[3] University of Liverpool, UK, email: wiebe@csc.liv.ac.uk
[4] University of Liverpool, UK, email: mjw@csc.liv.ac.uk

[5] Pauly [17] uses $[C]$ where we use $\langle C\rangle$.

[6] Strictly speaking, the structures of Coalition Logic associates a strategic game *form* to each state; a strategic game without preferences.

bring about both outcomes simultaneously. Finally, we do not want either agent to be able to unilaterally dictate an outcome – we want them both to have "equal power".

These requirements may be formally and naturally represented using CL, as follows:

$$\langle A, B\rangle p \quad (1)$$

$$\langle A, B\rangle q \quad (2)$$

$$\neg\langle A, B\rangle(p \wedge q) \quad (3)$$

$$\neg\langle A\rangle p \quad (4)$$

$$\neg\langle B\rangle p \quad (5)$$

$$\neg\langle A\rangle q \quad (6)$$

$$\neg\langle B\rangle q \quad (7)$$

Property (1) states that A and B can collectively choose p, while (2) states that they can choose q; (3) states that they cannot choose p and q simultaneously; and properties (4)–(7) state that neither agent can dictate an outcome.

3 Quantification

3.1 Quantified Coalition Logic

Expressing many interaction properties requires *quantification over agents and/or coalitions*. For example, consider the following *weak veto player* property [22]: "no coalition which does not have agent i as a member can make φ come about". This property can indeed be expressed in Coalition Logic as follows:

$$\bigwedge_{C \subseteq (Ag \setminus \{i\})} \neg\langle C\rangle\varphi$$

where Ag is the set of all agents in the system (the *grand coalition*). We thus use conjunction as a universal quantifier. The problem with this formulation is that it results in a formula that is exponentially long in the number of agents in the system. An obvious solution would be to extend CL with a first-order-style apparatus for quantifying over coalitions. In such a quantified CL, one might express the above by the following formula:

$$\forall C : ((C \subseteq Ag \setminus \{i\}) \rightarrow \neg\langle C\rangle\varphi)$$

However, adding quantification in such a naive way leads to undecidability over infinite domains (using basic quantificational set theory we can define arithmetic), and very high computational complexity even over finite domains. The question therefore arises whether we can add quantification to cooperation logics in such a way that we can express useful properties of cooperation in games *without* making the resulting logic too computationally complex to be of practical interest. In [2], we answered this question in the affirmative. We introduced *Quantified Coalition Logic (*QCL*)*, which allows a useful but restricted form of quantification over coalitions. In QCL, we replace cooperation modalities $\langle C\rangle$ with expressions $\langle P\rangle\phi$ and $[P]\phi$; here, P is a *predicate over coalitions*, and the two sentences express the facts that *there exists a coalition C satisfying property P such that C can achieve* ϕ and *all coalitions satisfying property P can achieve* ϕ, respectively. Examples of coalition predicates are, when C' is a coalition and n is a natural number:

- *supseteq*(C'): satisfied by a coalition C iff C is a superset of C'
- *geq*(n): satisfied by a coalition C iff C contains more than or equal to n agents
- *gt*(n): satisfied by a coalition C iff C contains more than n agents
- *maj*(n): satisfied by a coalition C iff C contains more than $n/2$ agents

For example, the property that agent i is a weak veto player for φ can be expressed as $\neg\langle\neg supseteq\{i\}\rangle\varphi$. Here the expression does not depend on the number of agents in the system. Thus we add a limited form of quantification *without* the apparatus of quantificational set theory. The resulting logic, QCL, is exponentially more succinct than the corresponding fragment of CL, while being computationally no worse with respect to the key problem of model checking.

To see how QCL makes it easier to express properties related to *voting*, consider the specification of *majority voting*:

> *An electorate of n voters wishes to select one of two outcomes* ω_1 *and* ω_2*. They want to use a simple majority voting protocol, so that outcome* ω_i *will be selected iff a majority of the n voters state a preference for it. No coalition of less than majority size should be able to select an outcome, and* any *majority should be able to choose the outcome (i.e., the selection procedure is not influenced by the "names" of the agents in a coalition).*

Let *maj*(n) be a predicate over coalitions that is satisfied if the coalition against which it is evaluated contains a majority of n agents. For example, if $n = 3$, then coalition $\{1, 3\}$ would satisfy the predicate, as would coalitions $\{2, 3\}$ and $\{1, 2\}$, but coalitions $\{1\}$, $\{2\}$, and $\{3\}$ would not. We can express the majority voting requirements above as follows. First: *every majority should be able to select an outcome*.

$$([maj(n)]\omega_1) \wedge ([maj(n)]\omega_2)$$

Second: no coalition that is not a majority can select an outcome.

$$(\neg\langle\neg maj(n)\rangle\omega_1) \wedge (\neg\langle\neg maj(n)\rangle\omega_2)$$

Simple though this example is, it is worth bearing in mind that its expression in CL is exponentially long in n.

3.2 Quantified Epistemic Logic

Epistemic logics [11, 14] give an account of agents' knowledge or beliefs. Operators K_i, C_G, E_G and D_G where i is an agent and G is a coalition are often used; $K_i\phi$, $C_G\phi$, $E_G\phi$ and $D_G\phi$ mean that i knows ϕ, that ϕ is common knowledge in the group G, that every member of G knows ϕ, and that ϕ is distributed knowledge in G, respectively.

The C_G, E_G and D_G operators let us express properties about group knowledge, but certain properties require quantification over agents and/or coalitions. Consider, for example, the following property:

> At least two agents know that at most three agents know ϕ, from an overall set of agents $\{1, 2, 3, 4\}$.

A way to express this fact in conventional epistemic logic is as follows:

$$\begin{aligned}&E_{\{1,2\}}\psi \vee E_{\{1,3\}}\psi \vee E_{\{1,4\}}\psi \vee \\ &E_{\{2,3\}}\psi \vee E_{\{2,4\}}\psi \vee E_{\{3,4\}}\psi \vee \\ &E_{\{1,2,3\}}\psi \vee E_{\{1,2,4\}}\psi \vee E_{\{1,3,4\}}\psi \vee \\ &E_{\{2,3,4\}}\psi \vee E_{\{1,2,3,4\}}\psi\end{aligned}$$

where ψ is:

$$(\neg K_1\varphi \vee \neg K_2\varphi \vee \neg K_3\varphi \vee \neg K_4\varphi)$$

The problem with this expression is similar to the problem with quantifying over coalitions in (standard) coalition logic discussed above: it is not very succinct, exponentially long in the number of agents in the system, and unrealistic for practical purposes. Again, we could add first-order style quantifiers, making it possible to express the property above as

$$\exists G : (|G| \geq 2) \wedge E_G\psi,$$

but this approach has the same disadvantages as discussed in the coalition logic case above.

But now we have a tool for limited quantification: coalition predicates. In [1] we introduce an epistemic logic with quantification over coalitions (ELQC), where the C_G, E_G and D_G operators are replaced by operators $\langle P\rangle_C$ and $[P]_C$, $\langle P\rangle_E$ and $[P]_E$, and $\langle P\rangle_D$ and $[P]_D$, respectively, where P is a coalition predicate. Now, $\langle P\rangle_C\phi$ means that *there exists a coalition G satisfying property P such that G have common knowledge of* ϕ, $[P]_C\phi$ means that *all coalitions G satisfying property P have common knowledge of* ϕ, and similarly for the two other kinds of group knowledge. The property discussed above can now be expressed as:

$$\langle geq(2)\rangle_E\neg\langle gt(3)\rangle_E\phi.$$

Possibly interesting properties of voting protocols include their knowledge dynamics. For example, when the winner of a voting protocol is announced, what does that tell an agent or a group of agents about the votes of other agents? ELQC can be used to reason about such properties. As an example, consider the following situation.

> A committee consisting of Ann, Bill, Cath and Dave, vote for who should be the leader of the committee (it is possible to vote for oneself). The winner is decided by majority voting (majority means at least three votes, if there is no majority there is no winner).

Consider first the situation before the winner is announced. Let proposition a mean that Ann wins, and una_a that Ann wins unanimously, and similarly for the other three agents. The following ELQC formula holds (no matter what the actual votes are):

$$\neg a \rightarrow \langle geq(2)\rangle_D\neg\langle geq(3)\rangle_E(\neg una_b \wedge \neg una_c \wedge \neg una_d).$$

The formula says that if Ann does not win, there is a group of at least two agents who distributively know that at most two agents know that neither Bill nor Cath nor Dave wins unanimously.

Consider next the situation when, after the secret voting, the winner is announced to be Ann. Let *Vote* be a set of atomic formulae, each denoting a complete vote (e.g., "Ann, Bill and Cath voted for Ann, Dave voted for himself"). The ELQC formula

$$\bigwedge_{vote \in Votes} (vote \rightarrow [supseteq(\emptyset)]_C\langle gt(1)\rangle_E vote)$$

denotes the fact that no matter what the actual vote is, in any coalition it is common knowledge that at most one agent knows the actual (complete) vote. This formula is true after the winner is announced to be Ann.

4 Logics for Coalitional Games

As mentioned in Section 2, there is a strong connection between Coalition Logic and non-cooperative games. As a result of the inherent differences between the class of non-cooperative on the one hand and the class of *coalitional*, or *cooperative*, games (as studied in coalitional game theory [15, Part IV]), on the other, the usefulness of standard Coalition Logic in reasoning about the latter type of games is limited. One of the main questions related to coalitional games is: "Which coalitions will form?", or "Which coalitions are stable?". Solution concepts such as the *core* have been proposed in coalitional game theory in order to capture the idea of rational participation in a coalition. In [4, 3], we develop two logics, *Coalitional Game Logic* and *Modal Coalitional Game Logic*, for reasoning about such games. Both logics keep the main syntactic construct of Coalition Logic, but the formulae are now interpreted in the context of a (single) coalitional game.

A *coalitional game* (without transferable payoff) is an $(m + 3)$-tuple [15, p.268]: $\Gamma = \langle Ag, \Omega, \sqsupseteq_1, \ldots, \sqsupseteq_m, V\rangle$ where , $\sqsupseteq_i \subseteq \Omega \times \Omega$ is a complete, reflexive, and transitive *preference relation*, for each agent $i \in Ag$.

In [4, 3] we discuss these logics in detail, including axiomatisation, expressiveness and computational complexity.

4.1 Coalitional Game Logic

The main construct of Coalitional Game Logic (CGL) [4] is the CL construct $\langle C\rangle\varphi$, again with the intended meaning that C can make φ come about. In addition, CGL has symbols for referring to particular outcomes, as well as formulae of the form $\omega \succeq_i \omega'$, where ω and ω' are outcomes, meaning that agent i weakly prefers ω over ω'.

As an example of a coalitional game property expressed in CGL, take the following: "outcome ω is in the core". The *core* of a coalitional game is the set of outcomes which can be chosen by the grand coalition such that no coalition can choose a different outcome which is strictly better for all the agents in the coalition:

$$CM(\omega) \equiv \langle Ag\rangle\omega \wedge \neg\left[\bigvee_{C\subseteq Ag}\bigvee_{\omega'\in\Omega}(\langle C\rangle\omega') \wedge \bigwedge_{i\in C}(\omega' \succ_i \omega)\right]$$

expresses the fact that ω is a member of the core. The formula $CNE \equiv \bigvee_{\omega\in\Omega} CM(\omega)$ will then mean that the core is non-empty.

4.2 Modal Coalitional Game Logic

While the main construct of Modal Coalitional Game Logic (MCGL) [3] still is the familiar $\langle C\rangle\varphi$; its interpretation is here radically different: the intended meaning is that coalition *C prefers* φ. The formulae are now interpreted in the context of an outcome in a coalitional game, and $\langle C\rangle\varphi$ is true if there is some other outcome which is (weakly) preferred by every agent in C where φ is true. Similar modalities were used by Harrenstein [12] in the context of non-cooperative games. The operator $\langle C^s\rangle$ is used to denote *strict* preference in the same way, and the duals $[C]$ and $[C^s]$ denote that the formula is true in *all* outcomes preferred over the current outcome by all agents in C. In addition, the language has an atomic symbol p_C for each coalition C, meaning that the current outcome can be chosen by C.

The fact that the current outcome is in the core can now be expressed as:

$$MCM \equiv p_{Ag} \wedge \bigwedge_{C\subseteq Ag}[C^s]\neg p_C$$

Comparing CGL and MCGL, interpreted over games with a finite set of outcomes the former is more expressive than the latter. However, observe that the expression $CM(\omega)$ quantifies by taking a disjunction over all outcomes. When the set of outcomes are infinite, this property cannot be expressed in CGL. In contrast, MCGL can express solution concepts such as core membership (expressed by the property *MCM*) and non-emptiness of the core also for games with infinitely many outcomes. Furthermore, MCGL can express many properties more *succinctly* than CGL: observe the difference between $CM(\omega)$ and *MCM*.

5 Logics for Aggregation of Preferences and Judgments

Preference aggregation – the combination of individuals' preference relations over some set of alternatives into a single social preference relation – has been studied in social choice theory for quite a while. The following is an example of three individuals' preferences over three alternatives a, b, c:

1	$a > b$	$b > c$	$a > c$
2	$a > b$	$c > b$	$c > a$
3	$b > a$	$b > c$	$c > a$
PaMV	$a > b$	$b > c$	$c > a$

The example also shows the result of pair-wise majority voting (PaMV), and serves as an illustration of Condorcet's voting paradox: the result of PaMV is not always a proper preference relation (in the example it is cyclic). Arguably the most well known result in social choice theory is Arrow's theorem [8], saying that if there are more then two alternatives then no aggregation function can have all of a certain collection of reasonable properties (non-dictatorship, independence of irrelevant alternatives, Pareto optimality).

We argued above that CL can be used to express properties of social choice mechanisms. However, neither CL nor any of the other logics we have mentioned so far are expressive enough to be used for reasoning about certain important properties related to aggregation. A logic which can express such properties would be useful for, e.g., specifying and verifying electronic voting mechanisms.

Another link between aggregation and logic is the emerging field of *judgment aggregation* within social choice. Judgment aggregation is concerned with combining individuals' judgments on a set of logically interconnected propositions into a set of collective judgments on the same propositions. An example, illustrating voting in a committee on propositions "the candidate is qualified" (p), "if the candidate is qualified he will get an offer" ($p \rightarrow q$) and "the candidate will get an offer" (q) (Y[es]/N[o]):

	p	**p → q**	**q**
1	*Y*	*Y*	*Y*
2	*Y*	*N*	*N*
3	*N*	*Y*	*N*
PrMV	*Y*	*Y*	*N*

The example also shows the result of proposition-wise majority voting (PrMV), and serves as an illustration of the so-called *discursive dilemma*: although positions of the individual voters all are logically consistent, the result of PrMV is not. The similarity between Condorcet's paradox and the discursive dilemma suggests a relationship between classical Arrowian preference aggregation and judgment aggregation – and, indeed, recent research [10] shows that the former is a special case of the latter.

*Judgment Aggregation Logic (*JAL*)* [5, 6] was developed specifically for expressing properties about judgment aggregation mechanisms. In consequence, it can also be used for classical preference aggregation. We use the logic to study the relationship between preference aggregation and judgment aggregation. Being tailor made for the purpose, it is much more expressive than CL when it comes to aggregation. The logic can express, e.g.:

- aggregation rules such as pair-wise and proposition-wise majority voting;
- properties of aggregation mechanisms such as non-dictatorship, independence of irrelevant alternatives and Pareto optimality; and
- important results such as the discursive paradox, Arrow's theorem and Condorcet's paradox.

A sound and complete axiomatisation is provided. This effectively gives us a proof theory for social choice and judgment aggregation. For example, that we can express Arrow's theorem in the logic on the one hand, and that we have a sound and complete proof system on the other, mean that we have a formal way to prove Arrow's theorem. Thus, the logic might be useful not only as a tool for specifying and verifying computational mechanisms, but also as a computational tool for social choice.

We give a taste of JAL in the context of preference aggregation. The language has to pairs of dual modalities. $\Diamond$ ($\Box$) quantifies over preference profiles (one preference relation for each agent), and $\blacklozenge$ ($\blacksquare$) over pairs of alternatives; the diamonds denoting existential and the boxes universal quantification. In addition, there is an atomic formula i for each agent i, which is interpreted in the context of a preference profile and a pair $\langle a, b\rangle$ of alternatives, meaning that agent i prefers the first alternative (a) over the second (b). Finally, there is an atomic formula σ, interpreted in the context of an aggregation function, a preference profile and a pair of alternatives $\langle a, b\rangle$, meaning that in the result of aggregating the preferences the element a is preferred over b. Formulae can be seen as properties of aggregation functions. For example, the formula

$$ND = \bigwedge_{i \in \Sigma} \Diamond \blacklozenge \neg(\sigma \leftrightarrow i) \tag{8}$$

is the non-dictatorship property: for every agent i there is a preference profile and a pair of alternatives $\langle a, b\rangle$ such that it is not the case that both i and the aggregation prefers a over b. Another example is Pareto-optimality:

$$UNA = \Box \blacksquare ((1 \wedge \cdots \wedge n) \rightarrow \sigma) \tag{9}$$

REFERENCES

[1] T. Ågotnes, W. van der Hoek, and M. Wooldridge, 'Quantifying over coalitions in epistemic logic', in *Proceedings of the Seventh International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, eds., Lin Padgham and David C. Parkes. IFAMAAS, (May 2008). To appear.

[2] Thomas Ågotnes, Wiebe van der Hoek, and Michael Wooldridge, 'Quantified coalition logic', in *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence (IJCAI 2007)*, ed., M. M. Veloso, pp. 1181–1186, California, (2007). AAAI Press.

[3] Thomas Ågotnes, Wiebe van der Hoek, and Michael Wooldridge. Reasoning about coalitional games, 2008. Manuscript.

[4] Thomas Ågotnes, Michael Wooldridge, and Wiebe van der Hoek, 'On the logic of coalitional games', in *Proceedings of the Fifth International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, eds., P. Stone and G. Weiss, pp. 153–160, Hakodate, Japan, (May 2006). ACM Press.

[5] Thomas Ågotnes, Michael Wooldridge, and Wiebe van der Hoek, 'Towards a logic of social welfare', in *Proceedings of The 7th Conference on Logic and the Foundations of Game and Decision Theory (LOFT)*, eds., Giacomo Bonanno, Wiebe van der Hoek, and Michael Wooldridge, pp. 1–10, (July 2006).

[6] Thomas Ågotnes, Michael Wooldridge, and Wiebe van der Hoek, 'Reasoning about judgment and preference aggregation', in *Proceedings of the Sixth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2007)*, eds., M. Huhns and O. Shehory, pp. 554–561. IFAMAAS, (May 2007).

[7] R. Alur, T. A. Henzinger, and O. Kupferman, 'Alternating-time temporal logic', *Journal of the ACM*, **49**(5), 672–713, (September 2002).

[8] K. J. Arrow, *Social Choice and Individual Values*, Wiley, 1951.

[9] *Handbook of Social Choice and Welfare Volume 1*, eds., K. J. Arrow, A. K. Sen, and K. Suzumura, Elsevier Science Publishers B.V.: Amsterdam, The Netherlands, 2002.

[10] Franz Dietrich and Christian List, 'Arrow's theorem in judgment aggregation', *Social Choice and Welfare*, (2006). Forthcoming.
[11] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi, *Reasoning About Knowledge*, The MIT Press, Cambridge, Massachusetts, 1995.
[12] P. Harrenstein, *Logic in Conflict*, Ph.D. dissertation, Utrecht University, 2004.
[13] S. Kraus, *Strategic Negotiation in Multiagent Environments*, The MIT Press: Cambridge, MA, 2001.
[14] J.-J. Ch. Meyer and W. van der Hoek, *Epistemic Logic for AI and Computer Science*, Cambridge University Press: Cambridge, England, 1995.
[15] M. J. Osborne and A. Rubinstein, *A Course in Game Theory*, The MIT Press: Cambridge, MA, 1994.
[16] M. Pauly, *Logic for Social Software*, Ph.D. dissertation, University of Amsterdam, 2001. ILLC Dissertation Series 2001-10.
[17] M. Pauly, 'A modal logic for coalitional power in games', *Journal of Logic and Computation*, **12**(1), 149–166, (2002).
[18] M. Pauly and M. Wooldridge, 'Logic for mechanism design — a manifesto', in *Proceedings of the 2003 Workshop on Game Theory and Decision Theory in Agent Systems (GTDT-2003)*, Melbourne, Australia, (2003).
[19] J. S. Rosenschein and G. Zlotkin, *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*, The MIT Press: Cambridge, MA, 1994.
[20] T. Sandholm, 'Distributed rational decision making', in *Multiagent Systems*, ed., G. Weiß, 201–258, The MIT Press: Cambridge, MA, (1999).
[21] M. Wooldridge, *An Introduction to Multiagent Systems*, John Wiley & Sons, 2002.
[22] M. Wooldridge and P. E. Dunne, 'On the computational complexity of qualitative coalitional games', *Artificial Intelligence*, **158**(1), 27–73, (2004).
[23] Michael Wooldridge, Thomas Ågotnes, Paul E. Dunne, and Wiebe van der Hoek, 'Logic for automated mechanism design – a progress report', in *Proceedings of the Twenty-Second Conference on Artificial Intelligence (AAAI 2007)*, ed., AAAI Press, Vancouver, Canada, (July 2007).