Multi-Agent Systems' Negotiation Protocols for Cyber-Physical Systems: Results from a Systematic Literature Review

Davide Calvaresi[1,2], Kevin Appoggetti[3], Luca Lustrissimini[3], Mauro Marinoni[1], Paolo Sernani[3], Aldo F. Dragoni[3] and Michael Schumacher[2]
[1]*Scuola Superiore Sant'Anna, Pisa, Italy*
[2]*University of Applied Sciences Western Switzerland, Sierre, Switzerland*
[3]*Università Politecnica delle Marche, Ancona, Italy*

Keywords: Software Engineering, Negotiation, Multi-Agent Systems, Cyber-Physical Systems, Real-Time Systems.

Abstract: Cyber Physical Systems (CPS) require a multitude of components interacting among themselves and with the users to perform automatic actions, usually under unpredictable or uncertain conditions. Multi-Agent Systems (MAS) have emerged over the years as one of the major technological paradigms regulating interactions and negotiations among autonomous entities running under heterogeneous conditions. As such, MAS have the potential to support CPS in implementing a highly reconfigurable distributed thinking. However, some gaps are still present between MAS' features and the strict requirements of CPS. The most relevant is the lack of reliability, which is mainly due to specific features characterizing negotiation protocols. This paper presents a systematic literature review of MAS negotiation protocols aiming at providing a comprehensive overview of their strengths and limitations, examining both the assumptions and requirements set during their development. While this work confirms the potential of MAS in regulating the interactions among CPS components, the findings also highlight the absence of real-time compliance in current negotiation protocols. Strongly characterizing CPS, the capability to face strict time constraints could bridge the gap between MAS and CPS.

1 INTRODUCTION

Cyber-Physical Systems (CPS) are deeply rooted in our daily living. Interconnected electronic devices of any size (from wearable to huge drivers) compose heterogeneous systems operating in various domains (e.g., manufacturing (Hsieh, 2002), zero-energy buildings, near-zero automotive fatalities (Rajkumar et al., 2010), telerehabilitation (Calvaresi et al., 2017b), and e-health (Calvaresi et al., 2014)). Scalable across time and space, with the ability to cope with a scenario's uncertainty, privacy concerns and security issues, CPS and MAS are transforming the humans' control of the physical world. Usually, these systems employ sensors to collect data from the real world, process them, and then provide feedback, either to other entities, or directly affecting (e.g., via actuators) the real world. Such systems are capable and responsible for both performing hard-coded and automatic actions and dealing with unpredictable or uncertain situations requiring "intelligent" actions. The distributed nature of such systems opens the horizon to a multitude of possible synergies. Interactions among entities of same or different systems represent a fascinating world, which has been largely investigated by the scientific community. However, new arising challenges have still to be faced.

On the one hand, according to Calvaresi et al (Calvaresi et al., 2017a), Multi-Agent Systems (MAS) is one of the most prominent and promising "approaches" supporting Internet of Things (IoT) technologies and CPS. The adoption of a multi-agent framework can facilitate the implementation of cooperative/competitive distributed thinking, robustness, reconfigurability, reusability (e.g., components capabilities, functionalities, knowledge), and a partial technology independence (smoother migration among different technologies) (Bellifemine et al., 2007; Calvaresi et al., 2016b). On the other hand, CPS require strict dependably, stringent safety and security policies, resources efficiency, and real-time guarantees (Rajkumar et al., 2010). For example, a safe use of personal devices (e.g., wearable blood-sugar/pressure devices), reliable and timely information delivery, bounded risks in receiving wrong information (in terms of content and timing), privacy guar-

Calvaresi, D., Appoggetti, K., Lustrissimi, L., Marinoni, M., Sernani, P., Dragoni, A. and Schumacher, M.
Multi-Agent Systems' Negotiation Protocols for Cyber-Physical Systems: Results from a Systematic Literature Review.
DOI: 10.5220/0006594802240235
In *Proceedings of the 10th International Conference on Agents and Artificial Intelligence (ICAART 2018) - Volume 1*, pages 224-235
ISBN: 978-989-758-275-2

antees and systems overall stable are features strictly required in safety-critical CPS.

Although the advantages provided by the adoption of MAS are remarkable, the full compliance with the requirements of CPS is not met yet (Calvaresi et al., 2017a). Uncertainty in the environment, security attacks, limitations in cyber models, and errors in physical devices make ensuring the overall system robustness, security, and safety, a critical challenge. The *distributed* decision-making process is crucial in the above-mentioned systems, and the *negotiation process* is essential for their success.

Contribution
To reach consensus or just interact, MAS need several negotiation protocols (standard and not). To better understand such contributions, this work performs a Systematic Literature Review (SLR) of the most relevant negotiation protocols proposed in the scientific literature addressing the following features:

(i) *assumptions* have been detailed to define the characteristics of environments and systems in which the negotiation processes are operating;

(ii) *requirements* have been presented and related to the assumptions to define which objectives and constraints have been set;

(iii) *Strengths*, and *limitations* collected by the primary studies have been elaborated to highlight achievements and still open challenges.

Elaborating and summarizing the evidence, the criteria presented in Section 3 have been generated and discussed. Finally, considering the *reliability* as the main requirement of safety-critical CPS, the negotiation's characteristics, constraints, and bounds have been formalized. The paper is organized as follows: Section 2 presents the review process and data collection, Section 3 organizes and describes the obtained results, Section 4 briefly discusses the obtained results in key CPS. Finally, Section 5 concludes the paper.

2 DATA COLLECTION AND REVIEW PROCESS: THE METHODOLOGY

Retrieving, selecting, and analyzing existing literature has more relevance if performed systematically. Hence, this paper adheres the procedure suggested by (Kitchenham et al., 2009) and adapted by (Calvaresi et al., 2016a). Such a methodology is composed of three stages (see Figure 1), and it is rigorous and reproducible [1].

Firstly, **Planning the review** defines steps and constraints. Such a phase elaborates a generic free-form question in *structured research questions* (SRQs) which characterize the pillars of the whole protocol. By doing so, the outcome will be reproducible, reliable, and comparable. The second stage, **Performing the review**, deals with the execution of the planned activities: *(i)* papers' collection and selection, *(ii)* paper elaboration, and *(iii)* features extraction. The last step, **Document Review**, deals with the data analysis and reporting activities related to the scientific dissemination.

2.1 Planning the Review

Defining the review process sets the research questions and their contexts, *search strategy*, *review protocol*, and *biases and disagreement resolution*.

Research Questions Definition
Investigating the scenarios presented in Section 1, the following free-form questions arose: *(i)* What *needs*, characterize the negotiations among agents in the several *application scenarios*? *(ii)* Are the solutions proposed by the scientific community satisfactory? *(iii)* How are such solutions characterized?

The Goal-Question-Metric (GQM), proposed by Kitchenham et al. (Kitchenham et al., 2010) and Galster et al. (Galster et al., 2014), ruled the decomposition of the unstructured questions mentioned above, into a set of three structured research questions. In particular, the *assumptions, requirements, strengths*, and *limitations* led the investigation and the definition of the following questions:

SRQ1 Setting the next question we aim at understanding the *Step 0* of the negotiation protocol development: *What are the **assumptions** rooting the most relevant approaches?*

SRQ2 To identify the goals targeted by such protocol, the following question is set: *What are the **requirements** such approaches intend to meet?*

SRQ3 The adoption of a specific negotiation algorithm would possibly bring some *advantages*. To name them, the following question is set: *What are the **strengths** and **limitations** characterizing the related negotiation approaches?*

Develop the Review Protocol
Once completed the definition of the structured-research-questions, the definition of the *Search Strategy* follows. Gray literature may introduce possible

[1] Primary studies selected and elaborated in early 2017

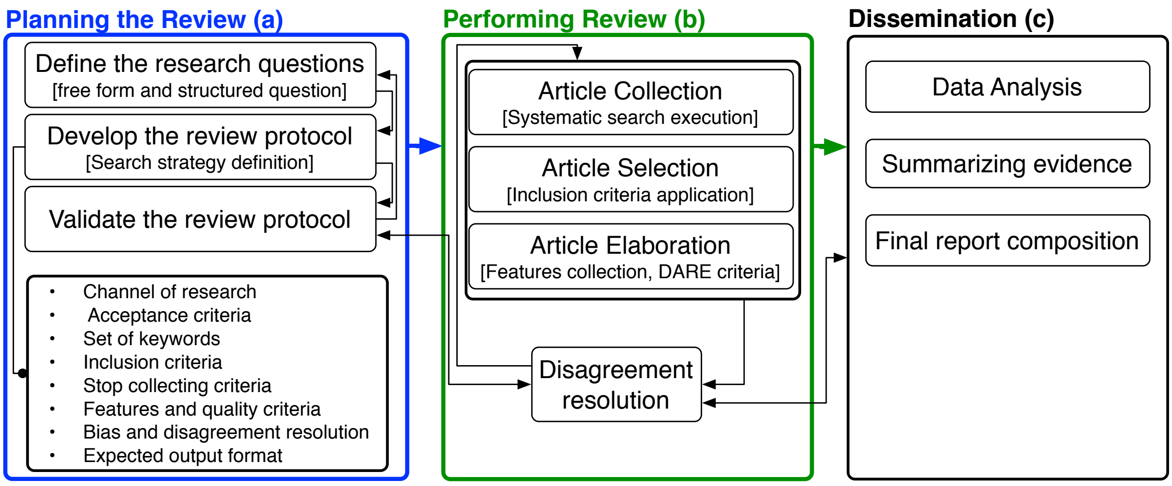

Figure 1: Review Methodology Structure according to (Kitchenham et al., 2009) and (Calvaresi et al., 2016a).

biases. Thus, only peer-reviewed collectors of papers (ieeeXplore[2], Sciencedirect[3], ACM Digital Library[4], and Citeseerx[5]) have been investigated.

To obtain more accurate results during the semi-automatic research, some keywords have been contextualized (by aggregating at least two or three words). According to the reviewers' rooted backgrounds and knowledge related to the Multi-Agent domain, the following set of keywords has been defined: *multi-agent interaction protocol, multi-agent negotiation protocol, agent-based negotiation, multi-agent problem-solving negotiation, distributed problem-solving negotiation, control distributed problem-solving*. For each query, the papers crawlers produced lists of articles ordered by pertinence. The criteria used to stop the paper collection is the same adopted by Calvaresi et al. in (Calvaresi et al., 2016a).

Inclusion Criteria Definition

The initial research counted 200 papers. A further coarse-grained examination reduced them to 143. The reviewers filtered them by performing a simultaneous and autonomous check of titles and abstracts' pertinence with the following *inclusion criteria*:

A) **Context:** The primary studies should define their contributions in the context of distributed-like systems;

B) **Purpose:** The purpose of primary studies should refer to mechanisms for negotiating tasks and resources or for achieving agreement or consensus between distributed entities.

[2]http://ieeexplore.ieee.org/Xplore/home.jsp
[3]http://www.sciencedirect.com/
[4]http://dl.acm.org/
[5]http://citeseerx.ist.psu.edu/index

C) **Relevance:** The primary studies should provide at least one of the following elements: [theoretical model, interaction mechanisms, practical implementation, tests, critical analysis, critical evaluations or discussion]

In the case of a clear verdict was missing (e.g., R1(*Yes*), R2(*No*), R3(*Maybe*)) the disagreement resolution process described below has been applied.

Features and Quality Criteria Definition

During the *"Features Collection"*, assessing the quality of the information provided by the primary studies is one of the main challenges of a Systematic Literature Review (Calvaresi et al., 2016a).

Although this work deals with a well-defined set of feature, context, rationale, research justification, critical examination, statement of findings and possible biases can hamper the credibility. Thus, the retrieved features have been classified by associating them Y - information is explicitly defined / evaluated, P - information is implicit / stated, or N - information is not inferable (DARE critirea (Kitchenham et al., 2009)).

Biases and Disagreement Resolution

The following expedients have been adopted to minimize and solve possible biases and conflicts. Developing the method and elaborating the articles, most of the tasks have been cross-checked. In particular, concerning Figure 1:

- the reviewers conducted the tasks included in 1(a) and (b) "*Planning the Review*", and "*Document Review*" collaborating synchronously.
- The collected articles list has been divided into three (number of reviewers performing the "*Article selection*") subsets, which have been pro-

cessed (applying the inclusion criteria check) by at least two out of three reviewers. The single reviewer's choices (*Yes, No, or Maybe*) have been kept hidden from each other till all of them had completed such a task. In the case of possible uncertainties (e.g., Yes-No, Yes-Maybe, No-Maybe) a third reviewer has been asked an extra check to finally decide weather include the article in the final list (to be elaborated) or not.

- During the "*Article Elaboration*", in the case relevant doubts arose, periodical collaborative *disagreement resolution meetings* have been organized.

3 RESULTS PRESENTATION

This section discusses the outcomes obtained by performing the methodology presented in Section 2. The main investigated issues are the assumptions on which the studied protocols rely on, the subsequent requirements set by the authors of the primary studies to identify and profile the proposed algorithms, and finally, the elaborated strengths and limitations, to summon the state of the art and identify future challenges.

3.1 Assumptions

The assumptions have been clustered to elicit abstract categories thus facilitating presentation and understanding (see Table 1). Most of the systems composed by distributed entities are based on the interactions among the available components. In MASs, such interactions have always been assumed asynchronous (Smith, 1980; Smith and Davis, 1981) strengthening the autonomy of single agents (e.g., their ability to execute without a direct human intervention and with full control over their own thread). Despite the communication-delay can be a crucial component, some studies neglect it, referring to the hypothesis of instantaneous message delivery (Aknine, 1998). In most cases, the authors refer to a general multi-agent architecture, even if few of the analyzed papers base their agents on the BDI paradigm (Atkinson et al., 2005). The design of a negotiation protocol mainly relies on the capability of taking autonomous decisions to pursue beliefs or directly self-interested or common goals. Indeed, the rationality (e.g., the ability of agents to always execute to achieve their goals, and never to prevent them from being achieved) and autonomy of agents are the most common assumptions in the analyzed studies. For example, in a group choice design support system (GCDSS), the agents negotiate on behalf of their user trying to persuade other agents according to their imposed or independently developed knowledge (Russell et al., 1995; Ito and Shintani, 1997).

Often, such autonomy has to face the impossibility of having agents ready with complete knowledge. Although dealing with partial knowledge might lead to possible deception, it is the most studied scenario in both cooperative and competitive MAS (Aknine et al., 2004; Zlotkin and Rosenschein, 1991; Smith and Davis, 1981). Having a competitive rather than cooperative agents' community, frames completely different scenarios and conditions which are even more complex in the case they are both cooperative and competitive at the same time. Some practical examples of negotiating limited knowledge in cooperative scenarios are the control of UAVs' task scheduling (Budaev et al., 2016), monitoring electricity transformation networks, and scheduling meetings (Kraus, 1997). Agents can collaborate by following self-organizing policies or relying on an orchestrator/coordinator (Wang et al., 2014) (the specular role in competitive scenarios is named "moderator" (Hanachi and Sibertin-Blanc, 2004)). Agents have to be "certified" or "trusted" (Alberti et al., 2004). Thus, the collaboration is more secure and can be applied in crucial activities such as decision making, coordination, and control processes. The bid-based negotiation approach is the most diffused, despite the involvement of simple or complex tasks (Aknine et al., 2004). In this approach, each agent can play two main roles: *(i)* the initiator (who calls for bids) and *(ii)* the contractor (who bids) in 1-to-1, 1-to-many scenarios, or auction based many-to-many (Wang et al., 2014). It can be predicted to last for short (Faratin et al., 1998) or long (Collins and Wolfgang Ketter, 2002) periods of time. In the scenario where the negotiation is still not converging, it might be considered as failed (Aknine et al., 2004). During a single instance of the bid-based protocols, an agent can play one of the two roles. Nevertheless, during the system execution, several negotiations of several tasks or resources can happen, and then, agents can play both *(i)* and *(ii)* (assuming a community of agents playing exclusively either *(i)* or *(ii)* is a rare scenario). In collaborative scenarios, due to their inner mechanisms, particular negotiation protocols need to prevent agents from over-bidding (e.g., very high rates in the Pre-Bidding phase). The solutions have been "bounding" the cooperation with the introduction of self-interested agents (Aknine et al., 2004), imposing "sequentiality" (Hanachi and Sibertin-Blanc, 2004), or limiting the number of issues to be possibly negotiated (Faratin et al., 1998).

Table 1: Assumptions overview.

Assumption	Class	Assumption	Class	Assumption	Class
No-commitment	AU	Agents multi-role	AR	Stationary/mobile agents	AR
Customizable neg.	FL	Customizable interaction prot.	FL	No comm-delay	AR
Mobile agents	AR	Autonomous agents	AU	agent as service provider	AR
Cooperative agents	CP	Neighborhood limited comm.	IN	Instantaneous messaging	IN
Partial knowledge	RA	Cooperative decision making	CP	Multi-crit. decision making	RA
Cooperative control	CP	BDI agents	AR	No bids on conflicting plans	RB
Cooperative computation	CP	Certified agents	RL	One bid per agent per time	FL
Multilateral comm.	IN	Limited resources	RA	Low-level comm. protocol	IN
Shared resources	AR	Delegation	FL	No explicit utility transfer	RA
Concurrent agents	CM	Castable (constraints/agents)	FL	Guaranteed resource alloc.	RL
Sequential neg.	RB	Feedback mechanism	RL	Information Completeness	AR
Roles re-definable	FL	Bounded behaviors	RA	Competitive agents	CM
Coordinated agents	AR	Stationary agents	AR	Self-interested agents	CM
1-to1 negotiation	IN	Static environment	AR	Indivisible resources	AR
1-to-n negotiation	IN	Tasks fully preemptable	RB	Loosely-coupled agents	AR
Time efficiency	PR	Fault-tolerance	RB	Asynchronous agents	AR
Failing negotiations	RB	simple tasks	AR	Sub-optimality	PR
Neg. topic related	AR	Ontology	IN	Timed negotiation	N
Long-time negotiation	FL	Agents' specific role	AR	Limited services/issues	FL
Indirect interactions	IN	Short-time negotiation	FL	Multi-negotiation	FL
Allowed counter-offers	FL	No preemption	AR	Independent tasks	AR
Penalized de-committing	RB	Extendable agents	FL	Rational Agents	AR
Uncertainty	RB				

Legend

AR	Architectural	FL	Flexibility	RA	Rationality	CM	Competition	RB	Robustness
AU	Autonomy	CP	Cooperation	IN	Interaction	PR	Performance	RL	Reliability

The "pool" of agents able to take part in a negotiation might be subject to some constraints. For example, it can be restricted by the concept of neighborhood (Olfati-Saber et al., 2007; Budaev et al., 2016) which can have completely different outcomes if considering stationary agents (e.g., agents which execute always in the same node of a network), mobile agents (e.g., agents able to migrate to different nodes at runtime), or hybrid scenarios (Ferber and Gutknecht, 1998; Wang et al., 2014). In (Aknine et al., 2004), the agent selection for a task execution is based on several factors such as the position of the agent in its environment and its capacity to process information.

Reza et al. (Olfati-Saber et al., 2007) give crucial importance to the agents' autonomy, especially in the presence of possible link/node failures unexpected time-delay and possible changes in the network topology. The assumption of having a system capable of operating as expected even in the case one or more failures happen is quite strong. However, several studies such as (Aknine et al., 2004) adopted it, facing scenarios where faults are most likely to happen. Several studies made assumption enforcing the flexibility, but hampering (in some cases impeding) the reliability. For example, the possibility of breaking a commitment (the promise made for a task execution in the bidding phase), with (Wu, 2008; Zhou et al., 2004) or without penalty, is not remotely allowed (Odell et al., 2001; Odell et al., 2000). Assuming the possibility of delegating tasks to other agents, it would boost flexibility and efficiency but limit reliability and rationality. The possibility of preempting tasks/behaviors is reasonable. However, assuming complete preemptability coupled with the absence of explicit deadlines, and allowing the possibility of failing negotiations, identical outcomes might be generated: multiple deadlines missing or direct starvation (Krothapalli and Deshmukh, 1999; Aknine et al., 2004). Sharing resources is a common practice to enhance system flexibility, bounded by their availability (Wellman and Wurman, 1998). Several protocols consider the customization of the negotiation interactions (Mazouzi et al., 2002) possible by also providing a pre-set personalization mechanism (Demazeau, 1995; Purvis et al., 2003). The agents' roles might be assumed static or dynamic (Wang et al., 2014; Faratin et al., 1998).

3.2 Requirements

Once the most common and relevant assumptions have been framed, the next step is to investigate the prevailing requirements set for negotiation protocols in MAS (see Table 2). The agents' interaction leading to the achievement of consensus and self/community goals captured the most concerns. Many contributions provide only negotiation-baselines, and thus require the implementation of generic/ad-hoc heuristics (Wanyama and Far, 2007). According to Mazouzi et al. (Mazouzi et al., 2002), being able to identify how and when to validate protocols, evaluate their success, and explain the relationships between agents, are outstanding requirements that must be considered. Nonetheless, deciding whom to interact with (e.g., agents with a higher reputation should have better bearing than others) and when initiating the interaction in certain scenarios is also crucial (Ramchurn et al., 2004).

On one hand, having an organized structure (Ferber and Gutknecht, 1998) and a flexible and automated agent community (Kraus, 1997) capable of achieving desired goals without affecting somebody else autonomy (Marzougui and Barkaoui, 2013) are the most common elements characterizing the environments in which the negotiation protocols have to operate in. On the other hand, having feasible, balanced, converging and preserved individual rationality and privacy are the most common elements that the protocols should present (Wellman and Wurman, 1998). For example, feasibility (basic assumption or requirement associating all the approaches) involves the need for setting functionalities such as check-and-validation of task assignment (Hsieh, 2002). Some approaches resulted in being extremely tailored on certain use-cases. Thus, they set very precise requirements to address a relatively broad multitude of goals. For example, the impossibility for the contractor to quit a task after having started it (Aknine et al., 2004), the non-retractability of bids, and the non-returnability of products (Guttman and Maes, 1998) are requirements set to foster reliability, especially in time-dependent solutions (Collins and Wolfgang Ketter, 2002). Moreover, although insufficient to fully provide real-time guarantees, some solutions seek for the respect of deadlines and schedulability guarantees (Shen and Norrie, 1998).

To enhance stability, some authors set the compliance with precedence and temporal constraints (Wanyama and Far, 2007). The time dependency has also been interpreted as the agents' capability of conceding more rapidly if the deadline approaches (Faratin et al., 1998). Regarding resources, they are assumed limited. Thus, setting a requirement regulating resources access and consumption regarding the agent community and their environment is mandatory. In trusted and collaborative environments, setting some policies is required to protect agents from exploiting each other (Faratin et al., 1998) and to discourage counter-speculations (Collins et al., 1998b). Other approaches to avoid security issues propose the requirement to specifically define *payment and permission mechanisms* (Collins et al., 1998b), transactions and market architectures (Collins et al., 1998a), mandatory penalty policies (e.g., non-penalization for new entrance and changing agents' identity (Ramchurn et al., 2004)), agent reputation update rate, and formal specification for processes validation (Mazouzi et al., 2002).

Regarding robustness, systems are required to either avoid failures or to keep working if they do occur below a certain threshold. One solution proposed in the primary studies is to supply information about the contractor during task execution (Ouelhadj et al., 2005). In particular, Collins et al. (Collins et al., 1998b) and Hsieh et al. (Hsieh, 2002) propose the requirement of a robust exception handler and a method to solve resource conflicts. *Architectural requirements* have been another important and recurrent element in the primary studies. For example, to overcome orchestration and autonomy limitations, a moderator could be compulsory (supporting community's fairness) (Hanachi and Sibertin-Blanc, 2004). Finally, to enhance or attain a certain performance, scenario-driven converging time and maximum execution time per task set are required (Vulkan and Jennings, 2000).

Despite the lack of critical analysis found in many scientific contributions (Calvaresi et al., 2016a), the analyzed papers have often proposed interesting clues. The more practical the proposed solutions are, the more detailed is the analysis of strengths and limitations. The mainly theoretical contributions presented a broad range of claims from the more explicit and easily understandable to the more ambitious and ambiguous. By looking at the big picture, common traits also associate entirely different approaches. Moreover, clustering strength allowed to define a sort of hierarchical relevance of the arisen categories. Due to space restrictions, the above-mentioned process will not be addressed in this paper. Nevertheless, such categories can be easily understood, since they reflect the structures of Section 3.3 and Section 3.4

3.3 Strengths

Table 3 collects all the features identified as "strengths" by the primary studies. Although feasibility is at the base of every process/protocol, it is not always guaranteed, and thus many studies consider it a "strength". Hence, it is not trivial having a converging negotiation protocol (Hanachi and Sibertin-Blanc, 2004; Matt et al., 2006) and guaranteeing that

Table 2: Requirements Overview.

Requirement	Class	Requirement	Class	Requirement	Class
Specifics formalization	VL	Protocols validation	VL	Planned maintenance	AR
Protocol evaluation	VL	Relating agents	RA	Tradeoff community/autonomy	AU
Agents goal isolation	RL	Entity-change discouraged	RL	(who/when)-heuristics for neg	IN
Promotion community join	FL	Agents' reputation balancing	FL	Interactions reputation-based	RL
Organized structures	AR	Automated agent	AU	Fake transactions penalization	RL
Privacy preservation	RL	Individual rationality	RA	Increasing GDSSs intelligence	PR
Efficiency	PR	Policies feasibility	RL	Services/gods non-returnability	AR
Increase of compatibility	FL	Convergence & equilibrium	PR	Reasonable Converging Time	PR
Ontology-based neg.	IN	Presence of moderators	AR	Manager operating in parallel	IN
Context-based interactions	IN	time-limited neg.	IN	Heterogeneous transactions	IN
Online tasks introduction	FL	Enanching inter-connections	AR	Complete agents' knowledge	RA
Interconnected managers	AR	Limited managers visibility	RA	Reasonable Execution Time	AR
Auction strictly-ruled	RL	Bids non-retractability	AR	Unbreakable commitment	RL
global goals	AR	Precedence constraints	AR	Time-dependent neg.	IN
Fault-tolerant neg.	RL	Resource-dependent neg.	IN	No-unbalanced exploitation	RL
Complex neg. contracts	AR	Anti-frauds control	RL	Energy-balancing heuristics	PR
Global social goals	IN	NO counter-speculations	RL	Secure resource supply	RL
Optimal neg.	PR	Payment mechanisms	IN	Enable rich-semantic language	IN
Enabled alliances	IN	Robust exception handling	RL	Multiple/Parallel neg.	FL
Scalability	AR	retro-compatibility	IN	Common time reference	AR
Competitiveness	CM	Shared knowledge	AR	Estimating due dates	FL
Shared policies	AR	Costs estimable	FL	Multiple providers per service	AR
Competitive negotiation	CM	Free community In/Out	FL	Heuristic-based bids	FL
Cooperative framework	CP	Optimized coordination	PR	High-level comm. lagns	IN
Norms taxonomy	VL	trust mechanism	RL	Mass customization	PR
Holonic dynamics	AR	Conflict resolution proc.	RL	Breakable contracts	RL
Deadlines respect	RL	Overview methods	RL	Comm-traffic reduction	PR

Legend

AR	Architectural	FL	Flexibility	RA	Rationality	CM	Competition	VL	Validation
CP	Cooperation	IN	Interaction	PR	Performance	RL	Reliability	AU	Autonomy

a *deal* can always be achieved (Faratin et al., 1998). Vice-versa, in the case of failures, detection and explanation of success/failure are possible (El Fallah-Seghrouchni et al., 1999). A possible way to avoid failures due to *computational intractability* is to negotiate throughout a centralized scheduling unit (Kanchanasevee et al., 1999). Seeking for effectiveness and efficiency, many analyzed solutions are extremely specialized and employable only in specific situations (Sun and Wu, 2009; Wu, 2008). Nevertheless, it is possible to mention cases that allow *language independence (El Fallah-Seghrouchni et al., 1999), context independence* (Cardoso and Bordini, 2016) and *protocol re-utilization* (Mazouzi et al., 2002), even in diametrically opposed scenarios (e.g., cooperative and competitive) (Sandholm, 1993). Some protocols can deal with uncertain environments, avoiding unexpected behaviors (Ito et al., 2008) and providing a high level of formalization (Kraus, 1997) (relatively flexible (Alberti et al., 2004)).

Moreover, having a controllable protocol size and a tractable complexity (Mazouzi et al., 2002) helps to enhance the system's stability (Olfati-Saber et al., 2007). Supporting agent autonomy (Hanachi and Sibertin-Blanc, 2004), one has to cope with a broad set of constraints. For example, they are radically different if the scenarios considered are firmly structured and automated (Wang et al., 2014) (hierarchical MAS (Wellman and Wurman, 1998)) or less structured, but considerably dynamic (e.g., the system just requires to observe juridical, common-sense, and behavioral laws (Wu, 2008), or admits rule re-definition on the fly (Purvis et al., 2003)). Finding an optimal trade-off between *completeness* (the capability of finding the optimal solution) (Ito et al., 2008) and the computational cost is always needed.

MAS are considered distributed by nature, thus guaranteeing low computational costs (Olfati-Saber et al., 2007; Collins and Wolfgang Ketter, 2002; Hong-tao and Kang, 2016; Golfarelli et al., 1997) is broadly recognized as a major strength. Concerning agent interactions, the overall performance of the community can be enhanced by shortening global negotiation processes (Aknine et al., 2004), avoiding infinite plan expansion for recursive plans (Cardoso and Bordini, 2016), generally reducing traffic (Smith, 1980), avoiding the broadcast of request messages to all the agents (Shen and Norrie, 1998),

Table 3: Strengths overview.

Strength	Class	Strength	Class
Deal always possible	RL	Convergence of conversation	RL
Success/failure detection	VL	NO computational intractability	RL
Improved efficiency/effectiveness	PR	Communication lang/tech-independent	AR
Allocations context-independent	FL	Protocol reuse	FL
Cooperative/competitive compliant	AR	Dealing with uncertain environments	FL
High-level of formalization	AR	Flexible specification	FL
Controllable protocol size	FL	Tractable complexity	PR
Stability	RL	Ensure autonomy	AU
Allows automated negotiation	AR	Hierarchical agents	AR
Juridical/common-sense compliance	RL	Rules changing on the fly	FL
Low computational costs	PR	Success/failure explanation	RA
Shorter global negotiation processes	PR	No diverging/recursive plans	RL
Net traffic reduced	PR	Avoidance of broadcasting requests	PR
Reduced negotiation rounds	PR	Dynamic task allocation	AR
Fast reaction to unpredictability	PR	Contract compliance verifiable	RL
Preventable neg. with blocked agents	RB	Tasks-sets atomically negotiable	AR
Better resource utilization	PR	Services description not required	RA
Multiple heuristics employable	AU	Possibile parallel negotiations	PR
De-commitment reduction	PR	Complex interactions observable	RA
Qualitative/quantitative analysis	VL	Conflict Resolution in Natural Language	IN
Trusted neg. sessions	RL	Increased task execution probability	PR

Legend

AR	Architectural	FL	Flexibility	RA	Rationality	CM	Competition	VL	Validation
RB	Robustness	IN	Interaction	PR	Performance	RL	Reliability	AU	Autonomy

and reducing rounds (Wanyama and Far, 2007) and messages-per-negotiation (Garcia et al., 2017). Enabling dynamic task allocation (Ouelhadj et al., 2005) is crucial. Thus, increasing the probability of task execution (Budaev et al., 2016) is highly appreciated. In terms of performance, the capacity of checking contract compliance (Vokřínek et al., 2007), and preventing negotiations with blocked agents (Aknine et al., 2004), can limit unpredictability (further reduced in (Budaev et al., 2016) by decreasing the reaction time to unpredictable events). Moreover, other relevant studies mentioned the capability of: negotiating sets of tasks considering them as *atomic bargaining items* (Sandholm, 1993), improving the resource utilization (Xueguang and Haigang, 2004), relaxing some constraints in "trusted" negotiation sessions (e.g., no need for services description) (Collins et al., 1998b), implementing different heuristics (Cardoso and Bordini, 2016), reducing the decommitment ratio, and paralleling the negotiation processes (Aknine et al., 2004).

Finally, some approaches permit to be evaluated by executing formal studies (El Fallah-Seghrouchni et al., 1999) such as qualitative and quantitative analysis (Mazouzi et al., 2002), and conflict resolution in natural Language (Demazeau, 1995).

3.4 Limitations

Gathering and analyzing the limitations have been the most challenging step of the whole review process. They emerge in three main ways: *related to the proposed solution* (often implicit and hidden between the lines), *to other approaches presented in the state of the art*, or *to specific solutions used as comparison terms*.

The data elaboration, performed to avoid duplicated elements and to simplify their understanding, added a considerable overhead in the elaboration process. Although several primary studies share the same limitations, more than a hundred different instances can be enumerated. The output of such aggregation is summarized in Table 4.

Sorted by relevance, only the most relevant per class are presented. The main limitation that affects some elaborated protocol is the possibility of ending up in a deadlock (Mazouzi et al., 2002; Aknine et al., 2004; Golfarelli et al., 1997) which can entail catastrophic consequences. In the case of short bidding windows, both initiators and contractors may lose opportunities. In the opposite scenario, with long bidding windows, the whole system might be congested, thus collecting a cascade of failures. Particularly for those protocols only suitable for single issue negotiation (Chang and Woo, 1994) or unable to handle

Table 4: Limitations Overview.

Limitation	Class	Limitation	Class
Risk deadlock	RL	Limited to single-issue neg.	PR
Limited to sequential neg.	PR	Risk of not reaching stability	RL
Single Point Of Failure	RL	Limited Knowledge access	IN
Impossibility of any-time tactics	PR	Statistic constraints and system's features	PR
High net-traffic	PR	Not scalable	FL
Additional Overhead neglected	PR	High computational cost	PR
Strictly domain-dependent	FL	Competitive scenarios neglected	AR
Semantic neglected	IN	Protocol limiting interactions	IN
Low efficiency	PF	Optimal distribution unreachable	PR
Conflicting sub-optimal allocations	RL	No dynamic rescheduling	PR
Bounded applicability (issues/agents/interactions)	PR	Dynamics Non-analyzable	RL
Feasibility non-observable	RL	Execution's correctness non-observable	RL
Risk of injection	RL	Risk of collusion	RL

Legend

AR	Architectural	FL	Flexibility	PR	Performance	IN	Interaction	RL	Reliability

parallel negotiations (Sandholm, 1993). This instability (Ito et al., 2008; Golfarelli et al., 1997) does not come alone. Hence, some approaches introduce single points of failure (Krothapalli and Deshmukh, 1999) such as the coordinator or moderator which can also be affected by a limited knowledge (Hanachi and Sibertin-Blanc, 2004; Vulkan and Jennings, 2000). In the *"Open-For-All environment"* (Vulkan and Jennings, 2000), there is a more pronounced incapability to apply tactics at any instant (Faratin et al., 1998), difficulties in defining/updating constraints and system features (Hanachi and Sibertin-Blanc, 2004; Jennings et al., 2001), an uncontrolled network traffic growth (Jennings et al., 2001; Faratin et al., 1998), expansion issues (Krothapalli and Deshmukh, 1999), and neglected additional overheads (Singh et al., 2010) (e.g., due to increasing computational costs (Ito et al., 2008; Wan et al., 2007)) hamper dramatically the systems' scalability. In terms of reusability, certain approaches present limited application domain (Krothapalli and Deshmukh, 1999; Aknine, 1998) (e.g., not considering competitive agents (Sandholm, 1993)). Low level and technologically committed approaches do not consider the semantic (Smith, 1980), thus concurring to generate interaction issues (Mazouzi et al., 2002; Jian, 2008). In term of performance, several studies refer to a general "low performance" (Krothapalli and Deshmukh, 1999; Ito et al., 2007), inefficiency (Ito and Shintani, 1997), and "non-optimality" (Vulkan and Jennings, 2000; Zhou et al., 2004). In particular, some approaches do not offer automatic mechanisms (Shen and Norrie, 1998) for task/resource runtime rescheduling. In same cases, scaling issues and agents (Wan et al., 2007) may arise problems as well (e.g., in (Ito et al., 2008), no more than two agents and seven issues can be properly handled). For example, in (Wellman and Wurman, 1998) there is a lack of in-depth analysis mechanisms, and in (Hsieh, 2002) checking the feasibility can be difficult or impossible(referred to cooperative communities). Finally, in terms of security, checking or enforcing the course of conversation is not always possible (Hanachi and Sibertin-Blanc, 2004). Some protocols leave the door open to possible injections, allowing *"strategic lying"* (tricking agents into believing the liars are trustworthy. Thus, they can exploit the unaware agents) (Ramchurn et al., 2004). Agents collusion is also a limitation and hence, a limited amount of mechanisms deal with "agent reputation" preventing such undesired circumstances (Ramchurn et al., 2004).

4 DISCUSSION

Exploiting the MAS' capability of negotiating in CPS represents a great potential, and it will be one of the main challenges for MAS in the upcoming years. According to Calvaresi et al. (Calvaresi et al., 2017a), MAS are still not ready to face strict timing constraints which strongly characterize the CPS. Nevertheless, many characteristics of the investigated negotiation protocols confirm such a potential. The agents in MAS can be seen as distributed nodes in CPS. Hence, they are assumed as autonomous, concurrent, coordinated, rational, multi-role, self-interested and loosely coupled. Computational and functional capabilities, communication (asynchronous), resources, and knowledge are considered limited. Resources can be shared, tasks in the system can be independent, architectures can be heterogeneous, and a mechanism for fault-tolerance has to be feasible. Sub-optimal resource allocations have to be reached in polynomial time. Unfortunately, some assumptions profoundly characterizing many negotiation protocols make them unable to cope with the requirements of CPS. In particular, in the presence of safety-critical CPS, assumptions such as *"no-commitment is required*, *the possibility of delegations*, and *only a vaguely defined time efficiency"* hamper the system reliability. In terms of requirements, *the impossibility to quit a running task*, *the non retractability of bidding*, *the possibility of using different agent heuristics*, *the desired guarantee of respecting deadlines* (for manufactured goods), and *the presence of precedence constraints*, go in the same direction of many CPS requirements.

Nevertheless, requirements such as *the introduction of a mediator mechanism to "simplify" the system dynamics*, *the possibility for the agents of changing their nature/identity*, and *unconstrained permission of agents to participate in multiple bids and tasks*, cannot be accepted. Strength is strongly subjected to the combination of requirements and assumptions. Thus, given such biases, anything inferred may result in inconsistent hypothesis. Instead, in the same situation, analyzing the limitations gives already important clues. The algorithms can be defined as inadequate to be employed in safety-critical CPS due to *the lack of commitment constraints*, *the difficulties in checking the feasibility*, *breaking contracts allowed by simply "paying" penalties*, *admission of a single point of failure*, and *impossibility of being scalable*.

5 CONCLUSIONS

This paper proposed an SLR applied to 143 primary studies to explore the *assumptions* standing behind the negotiation protocol in MAS and the *requirements* the different approaches set. Finally, *strengths* and *limitations* have been investigated to understand what has been done and what is still missing from the safety-critical CPS perspective.

The negotiation process in such systems involves smart nodes in distributed networks. The conventional decision-making processes performed in CPS are subject to more stringent constraints with respect to the ones characterizing traditional agent-based applications. The limitations presented in 3.4 and discussed in Section 4 depict a scenario in which the most relevant missing feature is the *reliability*.

Under the same assumption, bridging the gap between MAS and CPS (e.g., enabling the respect of strict timing constraints) can unveil new application scenarios in domestic, manufacturing, and healthcare domains. Finally, the analyzed techniques assume to operate in trusted environments. So far, if such a hypothesis is missing, the risk of injections and collusions is quite high. Hence, security challenges appeared to be still open, requiring to secure the systems at several levels.

Further work shall include the identification of the reliability of the primary objective, and the sets of assumptions and requirements that have to be redefined accordingly. Consequently, MAS would have to be purged from the inadequate components, which consist of several interventions in terms of theoretical contributions and practical development of new mechanisms. The proposed enhancements regard the agent local scheduler, and the communication middleware properly coupled with a new negotiation protocol based on concepts such as utilization factor and resource reservation (Calvaresi et al., 2017a).

REFERENCES

Aknine, S. (1998). Issues in cooperative systems: Extending the contract net protocol. In *Intelligent Control (ISIC), 1998. Held jointly with IEEE International Symposium on Computational Intelligence in Robotics and Automation (CIRA), Intelligent Systems and Semiotics (ISAS), Proceedings*, pages 582–587.

Aknine, S., Pinson, S., and Shakun, M. F. (2004). An extended multi-agent negotiation protocol. *Autonomous Agents and Multi-Agent Systems*.

Alberti, M., Daolio, D., Torroni, P., Gavanelli, M., Lamma, E., and Mello, P. (2004). Specification and verification of agent interaction protocols in a logic-based system. In *Proceedings of the 2004 ACM symposium on Applied computing*, pages 72–78. ACM.

Atkinson, K., Bench-Capon, T., and Mcburney, P. (2005). A dialogue game protocol for multi-agent argument over proposals for action. *Autonomous Agents and Multi-Agent Systems*, 11(2):153–171.

Bellifemine, F. L., Caire, G., and Greenwood, D. (2007). *Developing multi-agent systems with JADE*, volume 7. John Wiley & Sons.

Budaev, D., Amelin, K., Voschuk, G., Skobelev, P., and Amelina, N. (2016). Real-time task scheduling for multi-agent control system of uav's group based on network-centric technology. In *Control, Decision and Information Technologies (CoDIT), 2016 International Conference on*, pages 378–381. IEEE.

Calvaresi, D., Cesarini, D., Sernani, P., Marinoni, M., Dragoni, A., and Sturm, A. (2016a). Exploring the ambient assisted living domain: a systematic review. *Journal of Ambient Intelligence and Humanized Computing*, pages 1–19.

Calvaresi, D., Claudi, A., Dragoni, A., Yu, E., Accattoli, D., and Sernani, P. (2014). A goal-oriented requirements engineering approach for the ambient assisted living domain. In *Proceedings of the 7th International Conference on PErvasive Technologies Related to Assistive Environments*, PETRA '14, pages 20:1–20:4.

Calvaresi, D., Marinoni, M., Sturm, A., Schumacher, M., and Buttazzo, G. (2017a). The challenge of real-time multi-agent systems for enabling iot and cps. *in proceedings of IEEE/WIC/ACM International Conference on Web Intelligence (WI'17)*.

Calvaresi, D., Schumacher, M., Marinoni, M., Hilfiker, R., Dragoni, A., and Buttazzo, G. (2017b). Agent-based systems for telerehabilitation: strengths, limitations and future challenges. In *proceedings of X Workshop on Agents Applied in Health Care*.

Calvaresi, D., Sernani, P., Marinoni, M., Claudi, A., Balsini, A., Dragoni, A. F., and Buttazzo, G. (2016b). A framework based on real-time os and multi-agents for intelligent autonomous robot competitions. In *2016 11th IEEE Symposium on Industrial Embedded Systems (SIES)*, pages 1–10.

Cardoso, R. C. and Bordini, R. H. (2016). Allocating social goals using the contract net protocol in online multi-agent planning. In *Intelligent Systems (BRACIS), 2016 5th Brazilian Conference on*, pages 199–204. IEEE.

Chang, M. K. and Woo, C. (1994). A speech-act-based negotiation protocol: design, implementation, and test use. *ACM Transactions on Information Systems*.

Collins, J., Tsvetovat, M., Mobasher, B., and Gini, M. (1998a). Magnet: A multi-agent contracting system for plan execution. In *Proc. of SIGMAN*.

Collins, J. and Wolfgang Ketter, M. G. (2002). A multi-agent negotiation testbed for contracting tasks with temporal and precedence constraints. *International Journal of Electronic Commerce*, 7(1):35–57.

Collins, J., Youngdahl, B., Jamison, S., Mobasher, B., and Gini, M. (1998b). A market architecture for multi-agent contracting. In *Proceedings of the second international conference on Autonomous agents*, pages 285–292. ACM.

Demazeau, Y. (1995). From interactions to collective behaviour in agent-based systems. In *Proceedings of the 1st. European Conference on Cognitive Science.*

El Fallah-Seghrouchni, A., Haddad, S., and Mazouzi, H. (1999). Protocol engineering for multi-agent interaction. In *European Workshop on Modelling Autonomous Agents in a Multi-Agent World*, pages 89–101. Springer.

Faratin, P., Sierra, C., and Jennings, N. R. (1998). Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems*, 24:159–182.

Ferber, J. and Gutknecht, O. (1998). A meta-model for the analysis and design of organizations in multi-agent systems. In *Multi Agent Systems, 1998. Proceedings. International Conference on*, pages 128–135. IEEE.

Galster, M., Weyns, D., Tofan, D., Michalik, B., and Avgeriou, P. (2014). Variability in software systems-a systematic literature review. *IEEE Transactions on Software Engineering*, 40(3):282–306.

Garcia, E., Cao, Y., and Casbeer, D. W. (2017). Periodic event-triggered synchronization of linear multi-agent systems with communication delays. *IEEE Transactions on Automatic Control*, 62(1):366–371.

Golfarelli, M., Maio, D., and Rizzi, S. (1997). A task-swap negotiation protocol based on the contract net paradigm. *DEIS, Universitá di Bologna*.

Guttman, R. H. and Maes, P. (1998). Cooperative vs. competitive multi-agent negotiations in retail electronic commerce. In *International Workshop on Cooperative Information Agents*, pages 135–147. Springer.

Hanachi, C. and Sibertin-Blanc, C. (2004). Protocol moderators as active middle-agents in multi-agent systems. *Autonomous Agents and Multi-Agent Systems*, 8(2):131–164.

Hong-tao, L. and Kang, F.-j. (2016). Distributed task allocation modeling based on agent topology and protocol for collaborative system. *Optik-International Journal for Light and Electron Optics*, 127(19):7776–7781.

Hsieh, F.-S. (2002). Modeling and control of holonic manufacturing systems based on extended contract net protocol. In *American Control Conference, 2002. Proceedings of the 2002*, volume 6, pages 5037–5042.

Ito, T., Hattori, H., and Klein, M. (2007). Multi-issue negotiation protocol for agents: Exploring nonlinear utility spaces. In *IJCAI*, volume 7.

Ito, T., Klein, M., and Hattori, H. (2008). A multi-issue negotiation protocol among agents with nonlinear utility functions. *Multiagent and Grid Systems*.

Ito, T. and Shintani, T. (1997). Persuasion among agents: An approach to implementing a group decision support system based on multi-agent negotiation. In *International Joint Conference on Artificial Intelligence*, volume 15, pages 592–599.

Jennings, N. R., Faratin, P., Lomuscio, A. R., Parsons, S., Wooldridge, M., and Sierra, C. (2001). Automated negotiation: prospects, methods and challenges. *Group Decision and Negotiation*, pages 199–215.

Jian, L. (2008). An agent bilateral multi-issue alternate bidding negotiation protocol based on reinforcement learning and its application in e-commerce. In *Int. Symposium on Electronic Commerce and Security,*, pages 217–220.

Kanchanasevee, P., Biswas, G., Kawamura, K., and Tamura, S. (1999). *Contract-net based scheduling for holonic manufacturing systems*. PhD thesis.

Kitchenham, B., Brereton, P., Turner, M., Niazi, M., Linkman, S., Pretorius, R., and Budgen, D. (2010). Refining the systematic literature review process-two participant-observer case studies. *Empirical Software Engineering*, 15(6):618–653.

Kitchenham, B., Pearl Brereton, O., Budgen, D., Turner, M., Bailey, J., and Linkman, S. (2009). Systematic literature reviews in software engineering - a systematic literature review. *Information and Software Technology*, 51(1):7–15.

Kraus, S. (1997). Negotiation and cooperation in multi-agent environments. *Artificial intelligence*, 94(1-2):79–97.

Krothapalli, N. K. C. and Deshmukh, A. V. (1999). Design of negotiation protocols for multi-agent manufacturing systems. *International journal of production research*, 37(7):1601–1624.

Marzougui, B. and Barkaoui, K. (2013). Interaction protocols in multi-agent systems based on agent petri nets model. *Interaction*, 4(7):2013.

Matt, P.-A., Toni, F., and Dionysiou, D. (2006). The distributed negotiation of egalitarian resource allocations. In *Proceedings of the 1st international workshop on computational social choice (COMSOC06)*, pages 304–316.

Mazouzi, H., Seghrouchni, A. E. F., and Haddad, S. (2002). Open protocol design for complex interactions in multi-agent systems. In *Proc. of 1st. international joint conference on Autonomous agents and multiagent systems*.

Odell, J., Parunak, H. V. D., and Bauer, B. (2000). Extending uml for agents. *Ann Arbor*, 1001:48103.

Odell, J. J., Parunak, H. V. D., and Bauer, B. (2001). Representing agent interaction protocols in uml. In *Agent-oriented software engineering*.

Olfati-Saber, R., Fax, J. A., and Murray, R. M. (2007). Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE*, 95(1):215–233.

Ouelhadj, D., Garibaldi, J., MacLaren, J., Sakellariou, R., Krishnakumar, K., and Meisels, A. (2005). A multi-agent infrastructure and a service level agreement negotiation protocol for robust scheduling in grid computing. In *EGC*.

Purvis, M., Nowostawski, M., Oliveira, M., and Cranefield, S. (2003). Multi-agent interaction protocols of e-business. In *International Conference on Intelligent Agent Technology. IEEE/WIC*, pages 318–324. IEEE.

Rajkumar, R. R., Lee, I., Sha, L., and Stankovic, J. (2010). Cyber-physical systems: the next computing revolution. In *Proceedings of the 47th Design Automation Conference*, pages 731–736. ACM.

Ramchurn, S. D., Huynh, D., and Jennings, N. R. (2004). Trust in multi-agent systems. *The Knowledge Engineering Review*, 19(1):1–25.

Russell, S., Norvig, P., and Intelligence, A. (1995). A modern approach. *Artificial Intelligence. Prentice-Hall, Egnlewood Cliffs*, 25:27.

Sandholm, T. (1993). An implementation of the contract net protocol based on marginal cost calculations. In *AAAI*, volume 93, pages 256–262.

Shen, W. and Norrie, D. H. (1998). An agent-based approach for dynamic manufacturing scheduling. In *Proc. of Workshop on Agent-Based Manufacturing*.

Singh, A., Juneja, D., and Sharma, A. (2010). Introducing trust establishment protocol in contract net protocol. In *Advances in Computer Engineering (ACE), 2010 International Conference on*, pages 59–63. IEEE.

Smith, R. G. (1980). The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Transactions on computers*.

Smith, R. G. and Davis, R. (1981). Frameworks for cooperation in distributed problem solving. *IEEE Transactions on systems, man, and cybernetics*.

Sun, D. and Wu, J. (2009). Multi-agent coordination based on contract net protocol. In *Int. Symposium on Intelligent Ubiquitous Computing and Education*.

Vokřínek, J., Bíba, J., Hodík, J., Vybíhal, J., and Pěchouček, M. (2007). Competitive contract net protocol. *Theory and Practice of Computer Science*.

Vulkan, N. and Jennings, N. R. (2000). Efficient mechanisms for the supply of services in multi-agent environments. *Decision Support Systems*, 28(1):5–19.

Wan, W.-n., Zhang, J.-q., and Wang, M. (2007). A multi-agent negotiation protocol based on extended case based reasoning. In *IV International Conference on Fuzzy Systems and Knowledge Discovery*, volume 4, pages 462–467. IEEE.

Wang, G., Wong, T., and Wang, X. (2014). A hybrid multi-agent negotiation protocol supporting agent mobility in virtual enterprises. *Information Sciences*, 282.

Wanyama, T. and Far, B. H. (2007). A protocol for multi-agent negotiation in a group-choice decision making process. *Journal of Network and Computer Applications*, 30(3):1173–1195.

Wellman, M. P. and Wurman, P. R. (1998). Market-aware agents for a multiagent world. *Robotics and Autonomous Systems*, 24(3-4):115–125.

Wu, J. (2008). Contract net protocol for coordination in multi-agent system. In *Intelligent Information Technology Application, 2008. IITA'08. Second International Symposium on*, volume 2, pages 1052–1058.

Xueguang, C. and Haigang, S. (2004). Further extensions of fipa contract net protocol: threshold plus doa. In *Proceedings of the 2004 ACM symposium on Applied computing*, pages 45–51. ACM.

Zhou, P.-C., Hong, B.-R., Wang, Y.-H., and Zhou, T. (2004). Multi-agent cooperative pursuit based on extended contract net protocol. In *Machine Learning and Cybernetics, 2004. Proceedings of 2004 International Conference on*, volume 1, pages 169–173. IEEE.

Zlotkin, G. and Rosenschein, J. S. (1991). Incomplete information and deception in multi-agent negotiation. In *IJCAI*, volume 91, pages 225–231.