

COMPUTATIONAL LOGICS AND AGENTS*

[A ROADMAP OF CURRENT TECHNOLOGIES AND FUTURE TRENDS]

Michael Fisher¹, Rafael H. Bordini², Benjamin Hirsch³, Paolo Torroni⁴

1. Department of Computer Science, University of Liverpool, United Kingdom
m.fisher@csc.liv.ac.uk
2. Department of Computer Science, University of Durham, United Kingdom
r.bordini@durham.ac.uk
3. DAI-Labor, Technische Universität Berlin, Germany
benjamin.hirsch@dai-labor.de
4. Dipartimento di Informatica, Elettronica e Sistemistica, Università di Bologna, Italy
paolo.torroni@unibo.it

Abstract

The concept of an *agent* is increasingly used in contemporary software applications, particularly those involving the Internet, autonomous systems, or cooperation. However, with dependability and safety in mind, it is vital that the mechanisms for representing and implementing agents are clear and consistent. Hence there has been a strong research effort directed at using formal logic as the basis for agent descriptions and agent implementation. Such a logical basis not only presents the clarity and consistency required but also allows for important techniques such as logical verification to be applied.

We present a roadmap of research into the use of computational logic in agent-based systems and survey much of the recent work in these areas. Even though, with such a rapidly changing field, it is impossible to cover every development, we aim to give the reader sufficient background to understand the current research problems and potential future developments in this maturing area.

KEYWORDS: logic; autonomy; rational agents; multi-agent systems; agent specification; agent verification; logical implementation

1 Introduction

Representing and manipulating systems comprising large numbers of evolving and interacting computational entities is difficult. The concept of an *agent* was introduced as a useful abstraction for describing such computational entities, with the larger structures forming *multi-agent systems*. Representing these computational elements as agents and agent *organisations* provides a simple and intuitive metaphor for both individual and collective computation. However, it is important that software designers have appropriate and semantically clear mechanisms available for representing not only how individual agents adapt and evolve, but also how agents interact and combine to form new systems. Without the ability to represent and control agents, the practical development of complex multi-agent systems remain difficult, and agents will not be trusted for use in critical applications.

*Work supported by CoLogNet → <http://www.colognet.org>

In this document, we describe the use of *formal logic* as the basis for both determining and representing the behaviour of individual agents as well as agent organisations. The fact that formal logics have been used for this purpose is not surprising since they provide a framework that is both unambiguous and well understood, and in which specific logics can be designed to suit different scenarios. In particular, when an appropriate logic is chosen, it can provide a level of abstraction close to the key concepts of the software to be developed, as is the case with logics for multi-agent systems. We provide a survey of how computational logic (i.e., formal logic applicable to computing, for decision procedures, operational programming languages, etc.) can be used for agent-based systems. In particular, we will address:

- *agent description*, through logical specification and semantic techniques;
- *agent implementation*, through the execution of logical statements; and
- *agent analysis*, through logic-based verification techniques.

The link between agents and logic is important to both sides: logic-based approaches to agents are likely to grow in importance especially as verification becomes more necessary, while agents (particularly those that are required to *reason*) represent a key challenge for computational logic developments and application. Thus, we aim to describe the state of the art in key areas concerning logic and agents. In addition, we will also consider future trends in computing where computational logic approaches to agents ought to be relevant — this will serve to identify key future research issues and key requirements for wider use of these techniques.

2 Logic-Based Agent Description [Formal Specification]

Although we will not delve in detail into the “what is an agent” debate, we will briefly outline the variety of agents we consider. While an object can be seen essentially as a structure *encapsulating* behaviour (e.g., state, methods, ...), an agent is an *autonomous* object. Thus, the agent makes its own decisions about what to pursue. We are particularly concerned with *rational agents*, which can be seen as agents that make their decisions in a rational and explainable way. Thus, one of the key new aspects that the agent abstraction brings is that, for both design and verification, we need to consider not just what they do but *why* they do it. Since agents are autonomous, understanding why an agent chooses a course of action is vital. In [WJ95a], a number of additional aspects of agents are described, in particular their *pro-active*, *reactive* and *social* attributes. Although we do not see these as essential, most of the agents considered here do, indeed, incorporate such notions.

A number of logical theories of (rational) agency have been developed, such as the BDI [RG91, RG95] and KARO [vdHvLM93a, vL96] frameworks. One reason for using logic in agent-based systems is that a formal semantics comes (more or less) for free. From this it follows that the semantics of logic-based agents is strongly dependent on their underlying logic. The foundations of such agent description frameworks are usually represented as (often complex) non-classical logics. In addition to their use in agent theories, where the basic representation of agency and rationality is explored, these logics often form the basis of formal methods for agent systems.

The leading agent theories and formal methods in this area all share similar logical properties. In particular, the logics used bring together:

- an *informational* component, so as to be able to represent the agent’s beliefs or knowledge;
- a *motivational* component, in order to representing the agent’s desires or goals; and
- a *dynamic* component, allowing the representation of the agent’s activity.

Thus, the predominant approaches use relevant logical combinations: the BDI approach combines dynamic aspects via propositional temporal logic, informational aspects via modal logic KD45, and motivational aspects via modal logic KD; the KARO approach combines dynamic aspects via propositional dynamic logic, informational aspects via modal logics S5/KD45, and motivational aspects via modal logic KD. Unfortunately, these combinations can often be complex, both in terms of understandability and reasoning [FHM⁺95, BDF⁺01].

To give a flavour of the different approaches, we provide below brief overviews of some approaches to logic-based agent description.

- **BDI**

The BDI framework was developed by Rao and Georgeff over a number of years [RG91, RG95] and actually comprises two aspects: an agent theory and a practical agent architecture. The former is a well-known agent logic, or rather a framework for agent logics, based on Bratman's theory of practical reasoning [Bra87]. The latter is a much-copied architecture for building agents that are able to *reason* about their environment and *deliberate* about their actions.

The BDI theory describes agent behaviour using 'mentalistic' attributes, in particular **B**eliefs, **D**esires, and **I**ntentions, and uses these three components to form an agent's *mental state*. Belief corresponds to the information the agent has about its environment; desires (or goals) represent *options* available to the agent; and Intentions are the goals the agent has chosen to tackle immediately. Rao and Georgeff use branching time temporal logic [ES89] to incorporate the dynamic aspect (of action within an environment), and a Kripke structure [Kri63] to model the mentalistic notions. Thus, each Kripke world consists of a time tree in which the result of executing (or failing to execute) each action is represented.

The basic BDI architectural model, for example exemplified within both the Distributed Multi-Agent Reasoning System (dMARS) and the Procedural Reasoning System (PRS) [GL87, PRS01, RG95] can be summarized as follows. Essentially, such agents:

- observe the world and the agent's internal state, and update the event queue to reflect the events that have been observed;
- generate new possible desires (tasks), by finding plans whose trigger event matches an event in the event queue;
- from this set of matching plans, select one for execution (an intended means);
- push the intended means onto an existing or new intention stack, according to whether or not the event is a subgoal; and
- select an intention stack, take the topmost plan (the intended means), and execute the next step of this current plan — if the step is an action, perform it, otherwise, if it is a subgoal, post this subgoal on the event queue.

This approach has been extremely influential and is the basis for several agent systems and programming languages such as the PRS [GL87, PRS01], dMARS [RG95, dKLW98], and AgentSpeak(L) [Rao96].

- **KARO**

The foundations of KARO [vdHvLM93a, vL96] originate in the work of Moore on knowledge and action [Moo84]. Its name represents the main ingredients of the theory: **K**nowledge, **A**bilities, **R**esults, and **O**pportunity. The core logic consists of dynamic logic [Har84], used to model abilities and opportunities, and the well-known S5 logic [FHM⁺95], used to model knowledge. Abilities and opportunities are formalised as *potentials* using propositional dynamic logic and

logics of knowledge. Numerous extensions exist, augmenting the core language with formalisations for commitments, wishes, and goals [vLMvdH97, vdHvLM97, vLvdHM98, MvdHvL99, vdHMvS99]. Proof methods have been investigated, as has the influence of (non-)determinism [vdHvLM93b, MdBvE⁺01, HDS⁺01].

- **Temporal Logics**

While some approaches use *dynamic logics* to represent the underlying dynamic activity of an agent, many also use *temporal logics* [Eme90, Gol92] for this purpose. A straightforward use of discrete, linear temporal logics in specifying (and, as we will see later, verifying and executing) agents is given in [Woo92, Fis95, FW97, Fis05a]; this approach is in turn based upon work on temporal specifications of distributed and concurrent processes [BFG⁺96]. Much of this work uses a separated normal form (SNF) [Fis97b], in which temporal formulæ describe what the agent must do *now*, what it must do *next*, and what it guarantees to do at *some* point in the future. Temporal formulæ in SNF form can be *directly executed*.

However, just as there is a *wide* variety of temporal logics, so there are numerous different uses of temporal logics in agent-based systems. For example, approaches to temporal semantics of agent computation have been developed using techniques as disparate as Temporal Logic of Actions [MZW03], *dense* temporal logics [Fis96], and Temporal Z [RKJ04]. In addition, temporal techniques have been developed for multi-agent systems as well as autonomous agents; see, for example, [Sin96, JTW01]. In many cases the basic underlying temporal formalism has been extended with other logical dimensions such as those for knowledge, belief, goals, and motivations (for example [Fis05a]). Thus, this seemingly simple language is flexible enough to describe individual agents in terms of their dynamic behaviour, how their knowledge/information evolves, and how their goals/motivations evolve.

- **Deontic Logic**

Deontic logic is a form of modal logic where operators concern aspects such as *permission* and *obligation*. Deontic notions have been the basis of a number of agent descriptions, two of which are the IMPACT and BOID systems. IMPACT is a framework for coordinating (using hierarchical planning) legacy systems which is based on a logical foundation combining time, uncertainty, and deontic [ES98] operators (we will discuss IMPACT further later).

An approach to agent programming that extends the BDI concepts with deontic notions, particularly obligations, is the BOID architecture (**B**elief, **O**bligations, **I**ntentions, **D**esires) [BDH⁺01, BDH⁺02]. The argument here is that rational agents need to take into consideration norms and obligations that are set by the (agent) society [Dig99], and so must reason about them. The BOID architecture is an extension of the BDP logic [Tho00], which in turn is rooted in Reiter's default logic [Rei80]. It is based on conflict resolution for conditional beliefs, desires, obligations, and intentions. Instead of using modalities to model each attitude, the different notions are defined by sets of rules described by propositional logic. Default rules for each set are iteratively applied to the current theory. Whenever a conflict arises, a preference relation over the default rules is used to solve it.

- **Alternating-Time Temporal Logic**

Temporal logics [Eme90] have been studied for many years and form the basis of much of the work on verification of computational systems, as we will see later. Such logics come, traditionally, in two flavours: linear-time or branching-time. A third, more general, alternative is Alternating-Time Temporal Logics (ATL), introduced in [AHK97] and extended in [AHK02]. An interesting characteristic of this alternative to standard temporal logics is that one can quantify, selectively,

over paths that are seen as possible outcomes of game-theoretic interactions (to represent, for example, the possible interactions of an agent and its environment). In ATL, the individual sources of non-determinism (i.e., choices on “actions” that effect changes in the system or environment) can be referred to explicitly, hence the interest in the use of such logics for describing and reasoning about multi-agent systems. For example, Ryan and Schobbens [RS02] use ATL to define a refinement relation between agents, which is useful for agent verification.

A combined logic of potential interest for multi-agent systems, introduced by van der Hoek and Wooldridge, is ATEL [vdHW03]: an epistemic version of ATL which allows one to reason about what knowledge (or common knowledge, etc.) a coalition can effectively bring about or, indeed, what knowledge is required from agents so that they can effectively cooperate to bring about certain states of affairs. In combined logics, however, the semantics of interactions among the various modalities can often lead to various difficulties and complexities.

- **Situation Calculus**

The Situation Calculus [McCH69, LPR98b] is a first order language with some second order features that was originally developed to formalise dynamic aspects of the world, including concepts of ability and belief [LPR98a, PR99]. The logic itself is rather versatile, and has been expanded to include, among other things, concurrency [GLL00], time [PR93], non-monotonic reasoning [Bak91], and events [Lak99, Lin95].

Formulae in the situation calculus are usually interpreted over a set of axioms. A situation can be seen as a possible world history, i.e., a sequence of actions that led to the current state. Situations change by means of *fluents* (functions that are applied to situations). Fluents have pre- and post-conditions in order to minimise the amount of possible successor states. The situation calculus has been used to specify and implement agents, as we will see when the Golog programming language is mentioned later in this paper.

- **Z**

The formal specification language Z [Spi92] is based on set theory and first order predicate calculus. It extends the use of these languages by allowing an additional mathematical type known as *schema type*. Z schemas consist of a declarative part that declares variables and their types, and a predicate part that relates and constrains those variables.

A number of attempts have been made to create a theory in Z upon which agents can be based, with probably the most successful being those languages that provide Z specifications for objects [MD00, Smi00]. However, capturing an agent’s *motivations* as part of the core of Z appears problematic. Rather than creating an original agent theory based on Z, several existing agent theories and languages have been modelled using the Z specification language, for example AgentSpeak(L) [dL98] and dMARS [dKLW98].

- **Nonmonotonic Logics**

Nonmonotonic logics are used to formalise plausible reasoning [Mor99]. Reiter’s *default logic* [Rei80], which we have mentioned in relation with the BOID architecture, is a milestone in this area. Its key feature is the notion of *default rules*, used to model what typically happens given some assumptions and under certain conditions. Beside Reiter’s work, among nonmonotonic logics we count: Moore’s *autoepistemic logic* [Moo85], relating nonmonotonicity with rational agent beliefs; McCarthy’s logic of *circumscription* [McC86], formalizing nonmonotonicity by means of second-order axioms that limit the extension of certain “abnormal” predicates; and Alchourron et al.’s work on Belief Revision [AGM85]. Here, the authors propose a number of postulates, known in the literature as AGM postulates, capturing the properties that a belief revision operator should

satisfy in order to be considered rational. It is essentially a study of nonmonotonic reasoning from the dynamic point of view, and as such it has been considered very suitable to model rational agents. The general principle that should govern belief revision is indeed that of *minimal change* of information, while adapting to new input from the outside. Among the nonmonotonic forms of reasoning that are very relevant in agent research, we particularly identify *hypothetical reasoning* and *abduction*. Originally intended to model the reasoning process that leads from observations to plausible explanations (i.e. *hypotheses* [HW65]), abduction has been very much adopted as a general modelling and reasoning framework to extend incomplete belief or knowledge bases. This is very appealing in the agent domain, where knowledge incompleteness typically arises either because agents have to interact with other agents based on limited information, or where agents are situated in dynamic environments, subject to continuous change and evolution.

Because of the intrinsic complexity of such theories, the application of nonmonotonic reasoning in the agent context is as yet limited, with a few exceptions. For this reason, *belief revision for resource-bounded agents* has become an increasingly popular research area, leading to implementable logics for agent [AJL06, ABH⁺06]. On the other hand, logic programming and extensions have provided efficient techniques to handle many types of nonmonotonic and hypothetical reasoning (we elaborate on this later).

3 Logic-Based Agent Implementation [Programming Languages]

As we have seen in the previous section, there are many ways in which logic can be used to specify (and so provide semantics for) agent-based systems. However, we are also interested in *building* agent-based systems using techniques based upon logic. In other words, we are interested in the development and application of *computational logic* in this area. Computational logic has been very successful in the past, particularly through languages such as Prolog, in providing high-level abstractions for programming complex systems, while retaining a close link to formal semantics. The transparent formal semantics of logic-based approaches allows the possibility of applying formal verification techniques, thus enhancing the trustworthiness of the software.

In this section, we will survey languages that have the notions of agency and computational logic at their core. There is a wide variety of such languages [BDD⁺05]; while it is impossible to cover all of them, we aim to highlight a range of languages that use logical machinery to implement rational agents. It should be noted that this covers much more than logic programming, which is but one approach to implementation using logic. For example, languages such as AGENT-0 and Go! contain some elements of functional programming, while METATEM is based on forward chaining model construction. In addition, while there is a wide range of logic-based languages that implement very basic forms of agency, these are often little more than extensions of logic programming that capture communication and encapsulation. Thus, we will concentrate on a number of different mechanisms for using computational logic to implement specifically *rational* agents.

- **AGENT-0**

One of the seminal works in this area was produced by Shoham [Sho93] who proposed *Agent-Oriented Programming* (AOP) as a new paradigm for implementing (rational) agents. This framework consists of a logical system, called AGENT-0, for defining the mental state of agents, together with an interpreted language for programming those agents [Tho95]. The logical component of AGENT-0 is a quantified multi-modal logic, with three modalities: belief, commitment, and ability. An example of an acceptable formula of AGENT-0 is as follows:

$$CAN_{pele}^4 pass(ball)^9 \Rightarrow B_{maradona}^4 CAN_{maradona}^{10} score(goal)^{11}$$

The intended meaning of this statements is

if at time 4 pele can ensure that he will pass the ball at time 9, then at time 4 maradona believes that at time 10 he (maradona) can ensure a goal is scored at time 11.

As can be seen from the example above, the logical basis of AOP contains notions of time, as well as other modalities, and the programming language provides syntactic support for such concepts (though it is not clear what the formal link between the programming language and the logic is [WJ95a]). Rather than being widely used itself, it is the idea of a language for developing agents based upon multi-modal logics that has been taken up by other researchers. AGENT-0 has been the inspiration for many subsequent agent programming languages.

- **METATEM and Concurrent METATEM**

While most agent specifications are based on a combination of temporal or dynamic and modal logics, METATEM [BFG⁺95] was initially based on *temporal logic* alone. Temporal formulæ in a particular normal form (called SNF [Fis97b]) describe what the agent must do *now*, what it must do *next*, and what it guarantees to do at *some* point in the future. Most importantly, temporal formulæ in SNF form can be *directly executed* in order to implement the behaviour of an agent. This basis was extended for increasingly sophisticated forms of agency [Fis95], involving belief, deliberation [Fis97a], resource-bounded reasoning [FG99] and confidence [FG02]. In parallel, the language was developed into a concurrent version, for multi-agent systems [Fis94, Fis05b] and, more recently, work on this has moved towards flexible group organisation [FGH03], multi-agent collaboration [FGH04], and applications in pervasive/ubiquitous computing [HFGB05]. For an overview of these developments, see [Fis05b].

- **Golog and ConGolog**

Golog [HRL⁺97] is a high-level specification language, based on the situation calculus, extending it with control structures such as sequencing, test actions, non-deterministic choice (of both actions and programs), and procedures (essentially macros). The execution of a Golog program involves proving the theorem expressed by the program with respect to some background theory (expressed by its axioms). As a byproduct, variables are instantiated and so provide an execution trace. In this, it is reminiscent of Prolog and, not surprisingly, a Golog interpreter can easily be described in Prolog.

As with the situation calculus, the basic Golog programming language can be extended considerably, for example with sensing actions relevant to robotics [Lak99]. One of the most prominent extensions is ConGolog, a concurrent variety of Golog [GLL00]. The concurrency described is not true concurrency but rather an interleaving of processes. (Of course, on a one-processor machine this is always the case, but “truly” concurrent programs let the operating system take care of the actual interleaving, setting only their priority.) However, since ConGolog incorporates thread priorities and interrupts, either coarse-grained or fine-grained interleaving can be described. In the extreme case, it is possible to interleave whole programs rather than program steps.

Golog has been particularly applied in the (cognitive) robotics area, but is also being developed for use in the semantic web [MS01, MS02].

- **AgentSpeak**

While BDI architectures have been used in many practical agent systems, there are few high-level programming languages based systematically on the principles of BDI theory. AgentSpeak(L) is such a language, defined to be an extension of logic programming incorporating beliefs, events, goals, actions, plans, intentions [Rao96]; its reasoning cycle is very similar to the dMARS cycle described above. The behaviour of the agent (i.e., its interaction with the environment) is dictated by the plans written in AgentSpeak. Although the beliefs, desires, and intentions of the agent are

not explicitly represented as modal formulæ, such mental attitudes can be formally ascribed or checked given a representation of an agent's current state.

The practical use of AgentSpeak is being facilitated by a new interpreter and multi-agent platform called *Jason* [Bor06]. The interpreter implements, in Java, the operational semantics of an extended version of AgentSpeak [BHV05]. *Jason* is available *open source* (see <http://jason.sourceforge.net>).

- **Logic Programming and Non-Monotonic Reasoning**

Following Kowalski and Sadri's work on the *observe-think-act* agent cycle and architecture to reconcile rationality with reactivity [KS99], a conspicuous body of literature introduced agent theories based on Logic Programming [Llo87]. This choice is often justified by the aim to link declarative semantic agent descriptions to execution models. Therefore, even when not as expressive as other formalisms, logic programming languages and variations become appealing when they come together with formal results that link declarative and operation semantics.

Among the various logic programming frameworks and languages, the most commonly used for agent theories are: Answer Set Programming [Bar03], Abductive Logic Programming [KKT93] for hypothetical and temporal reasoning, Dynamic Logic Programming [ALP⁺98] for knowledge and program updates, and Logic Programming without Negation as Failure [DK95] for agent deliberation.

- **3APL** [<http://www.cs.uu.nl/3apl>]

3APL is a programming language for implementing cognitive agents [HdBdHM98]. It provides programming constructs for implementing an agent's beliefs, goals, basic capabilities such as belief updates or motor actions, and a set of practical reasoning rules through which the agent's goals can be updated or revised [DdBDM03]. 3APL programs are executed by an interpreter that implements an deliberation cycle using such constructs.

The 3APL language is a combination of imperative and logic programming. From the imperative programming viewpoint, 3APL inherits the full range of regular programming constructs, including recursive procedures and state-based computation. States of agents in 3APL, however, are belief (or knowledge) bases, which are different from the usual variable assignments of imperative programming. From the computational logic perspective, answers to belief-base queries of a 3APL agent are proofs in the logic programming sense.

The authors of 3APL have also examined a number of related agent languages, including **Dribble** [RvHM03] and **GOAL** [HdBHM01].

- **IMPACT** [<http://www.cs.umd.edu/projects/impact>]

IMPACT [AOR⁺99] evolved through an international research project led by the University of Maryland. The main aim of the IMPACT project was to “develop both a theory as well as a software implementation facilitating the creation, deployment, interaction, and collaborative aspects of software agents in a heterogeneous, distributed environment” [ESP99, ES99, ESR00, EMS02]. Although IMPACT is only partially based on logic, providing a “thin” logic component on top of traditional data structures, the logic component involves time, uncertainty and deontic logics [ES98]. It also makes use of stable models [GL88] for implementing agent rationality.

On the practical side, IMPACT provides a set of servers (yellow pages, thesaurus, registration, type and interface) that facilitate agent interoperability in an application independent manner. It also provides an Agent Development Environment for creating, testing, and deploying agents. Thus, IMPACT provides techniques and tools to build agents on top of existing legacy code, and

algorithms supporting a variety of applications (command and control, logistics, e-commerce, data mining, etc.).

- **JIAC IV** [<http://www.jiac.de>]

JIAC is an agent based service-ware framework geared towards telematics and telecommunication services. It has three major elements: it employs a flexible component structure which allows agents to quickly create and dynamically adapt their capabilities; it is knowledge based, as the agent architecture follows a BDI approach, based on three-valued predicate logic; and its communication model is structured around the service concept. Agents are programmed using JADL (JIAC Agent Description Language), which comprises ontology definitions, declarative plan elements, and procedural script-like elements. Java objects can be invoked from JADL using Agent-Beans. Due to its strong connection with industry the framework features advanced security as well as management facilities.

- **SOCS** [<http://lia.deis.unibo.it/research/socs>]

SOCS (standing for Societies Of Computees) was an European project aimed at defining a computational logic model for the description, analysis and verification of global and open societies of heterogeneous *computees*, i.e., agents in computational logic [Ton06]. The main outcomes of the project are definitions of the agent and society models and the proof-procedures implementing the operational models. The computee model is based on the \mathcal{KGP} (*K*nowledge, *G*oals, *P*lans) agent architecture [KMS⁺04, ST06, Sad06], whereas the *SOCS* social model is based on Social Integrity Constraints (\mathcal{IC}_S) [ACG⁺04b, CGA⁺06]. The proof-procedures to implement the operational models are the *SCIFF* [AGL⁺05] for interaction verification and *CIFF* [EMS⁺04] for temporal reasoning and planning, both of them extensions of Fung and Kowalski's *IFF* abductive proof-procedure [FK97], and Gorgias [KM03] for preference reasoning and deliberation, based on LPwNF [DK95]. A prototype of *societies of computees* is implemented in the PROSOCS architecture [BED⁺06, ACG⁺06].

The distinguishing features of the \mathcal{KGP} model are: a modular architecture, where reasoning capabilities can be implemented independently of each other, and a flexible agent cycle that can be programmed using preference rules. The *SOCS* social model is mainly oriented towards specification and verification of agent interaction, and \mathcal{IC}_S , its main component, can be used to express commitment-based semantics for agent communication languages and to define agent protocols, as well as for run-time verification of whether agents comply such specifications.

- **Agent reasoning and coordination based on Abductive Logic Programming**

Kowalski *et al.* were the first authors to propose an agent architecture based on Abductive Logic Programming [KKT93], to reconcile reactivity with rationality. KS agents adopt an *observe-think-act* cycle, where *think* is a (fixed) number of computational steps that the IFF proof-procedure [FK97] uses to compute actions to perform, based on the agent's current beliefs, goals, and inputs. The agent goals are a conjunction of Prolog terms and integrity constraints, i.e., forward rules that must be true at all times during the operation of the agent.

The KS agent architecture has been extended by Dell'Acqua and colleagues [DNP02] to accommodate *updates*, i.e., dynamic changes to the agent programs. KS agents have also been used by Sadri *et al.* to implement agents that engage in negotiation dialogues in the context of resource allocation [STT01, STT02b, STT03]. A more elaborate agent model which can be considered as a descendant of KS agents is the \mathcal{KGP} model mentioned above.

Agent deliberation is the main focus of work done by Kakas *et al.* based on Logic Programming without Negation as Failure (LPwNF) [KM02, KM03, KTD04]. The agent decision making process is specified by a number of preference rules structured in a multi-layered program. LPwNF is

used to implement an argumentation framework where policies are combined with context rules. Argumentation is used to solve conflicts arising, for instance, from an agent policy giving directives that are incompatible with those given by the context rules.

Orthogonal interest has been devoted by several authors to the issue of coordination of multi-agent reasoning activity. Two noteworthy bodies of work are the research conducted by Ciampolini *et al.* on ALIAS [CLM⁺03, CLMT02] and that by Satoh *et al.* on speculative computation [SIIS00, Sat02, SY02]. ALIAS is a multi-agent architecture where agents can reason hypothetically, making assumptions about missing pieces of information (as for KS agents), and where they can form groups (“bunches”) to coordinate their reasoning activity so as to ensure that their beliefs about unknown information are consistent with one another’s integrity constraints. This kind of interaction is kept at a separate level from that of reasoning, and is described by a language called LAILA [CLMT02]. Speculative computation is another form of collective and coordinated hypothetical reasoning, based on abductive logic programming. Its main purpose is to enable agents to continue their execution while waiting for pieces of information from the outside, based on default assumptions, and to switch to a different operation in case some incoming information contradicts their defaults.

- **Dynamic Logic Programming and Updates**

An increasingly popular logic programming based formalism for modelling multi-agent reasoning is Answer Set Programming, and extensions of Dynamic Logic Programming to accommodate adaptation and changes. Leite *et al.* developed an agent architecture, called *MINERVA*, based on Multi-Dimensional Dynamic Logic Programming (*MDLP*) [LAP01] and on the update command language *LUPS* [APP⁺99], where several specialised, possibly concurrent, sub-agents perform various tasks while reading and manipulating a common knowledge base. Related to *MINERVA* is the work on DALI [CT02], an agent programming language encompassing basic patterns for reactivity, pro-activity, internal thinking, and memory. The semantics of a DALI program is defined in terms of another program, where reactive and proactive rules are reinterpreted in terms of standard Horn Clause rules.

- **April and Go!**

April [McC95] combines logic and functional programming. Although a high-level language, it offers a simple interface to native code in languages such as C. As the authors put it: “April is more an object based concurrent language with objects as processes”. Go! [McC04] is a descendant of April that is multi-threaded, strongly typed, and higher order. Agents in Go! contain both reactive and deliberative aspects, and coordinate using BDI structures.

- **Qu-Prolog**

Finally, Qu-Prolog [CRH01] is a multi-threaded and distributed version of Prolog with support for high-level communication between the threads. Agents can communicate using thread-to-thread messages, a publish-and-subscribe event notification system, or via the shared Prolog dynamic database. The first two alternatives allow threads in different Qu-Prolog processes anywhere on the internet to communicate; the last communication alternative is restricted to threads in the same process [CRZ06].

Agent Communication Languages and Ontologies

An important aspect of multi-agent systems development is agent interaction. For this, Agent Communication Languages (ACL) with appropriate formal semantics are required, as is the use of ontologies so that agents agree on the semantics of the vocabulary used in communication. In this section, we briefly mention the most prominent work on ACL and ontologies, discussing also their logical basis.

- **KQML** [<http://www.cs.umbc.edu/kqml>]

The Knowledge Query and Manipulation Language (KQML) is a language and protocol for knowledge exchange [LF94]. It was developed as part of a larger effort, the ARPA Knowledge Sharing Effort which was aimed at developing techniques and methodologies for building large-scale, sharable, reusable knowledge bases. KQML is both a message format and a message-handling protocol to support run-time knowledge sharing among agents. It can be used as a language for an application program to interact with an intelligent system or for two or more intelligent systems to share knowledge in support of cooperative problem solving.

KQML focuses on an extensible set of performatives, which define the permissible operations that agents may attempt on each other's knowledge and goal stores. The performatives comprise a substrate on which to develop higher-level models of inter-agent interaction such as the contract net protocol and negotiation techniques. In addition, KQML provides a basic architecture for knowledge sharing through a special class of agents called "communication facilitators", which coordinate the interactions of other agents. The ideas underlying the design of KQML have been explored through experimental prototype systems which were used to support several testbeds in areas such as concurrent engineering, intelligent design, and intelligent planning and scheduling. Even though FIPA (see below) is becoming the standard ACL, it took most of its ideas from KQML, which also has a historical importance in being the first ACL to achieve widespread use.

- **FIPA** [<http://www.fipa.org/activities/semantics.html>]

The Foundation for Intelligent Physical Agents (FIPA) was formed in 1996 to produce software standards for heterogeneous and interacting agents and agent-based systems. In the development of these standards, FIPA obtained input and collaboration from its membership and from the agent's field in general to build specifications that can be used to achieve interoperability between agent-based systems developed by different companies and organisations.

FIPA is currently working on a new semantic framework to reflect the needs of verifiability and conformance. In particular, the objective is to adopt or define a semantic framework that can give an account of FIPA's existing aspects as well as a number of additional constructs such as contracts, agreements, policies, and trust.

- **OWL** [<http://www.w3.org/2004/OWL/>]

It is increasingly clear and widely recognised that in order to provide true interoperability, it is not enough to standardise on agent communication languages. The actual content of the messages also has to be interpreted in a semantically compatible way. Interpretation of information by computing systems is at the centre of the semantic web vision. In that area, the aim is to augment the syntactic description of web pages with semantic information.

While HTML allows us to display the information available on the web, it does not provide much capability for describing the information in ways that facilitate the use of software programs to find or interpret it. As an example, the word 'agent' on a web site might represent an aspect of Computer Science, an aspect of Chemistry, or a participant of an espionage plot.

OWL [McGvH04], the successor of the DAML+OIL [W3C1, HP-S02] language, provides a rich set of constructs with which to create ontologies and to markup information so that it is machine readable and understandable. A key approach here is the use of Description Logics [BCM⁺03], which allows reasoning about semantic information. There are three variants of the OWL language: OWL-light, OWL-DL (description logic), and OWL-Full. While OWL-Light allows for easy and efficient implementations, OWL-DL makes use of the expressive power of description logic while still being computable. OWL-Full violates some constraints of description logic [OWL]. OWL is a W3C Recommendation since February 2004.

4 Logic-Based Agent Verification [Agent Analysis]

Agents appear to be an appropriate abstraction for many systems that operate in complex, dynamic, and unpredictable environments. These include application areas such as air-traffic control, autonomous spacecraft control, health care, and industrial systems control, to name just a few. These are areas of application that clearly require *dependable* computational systems and, in particular, that such systems should be *verifiable*. In this section we address this issue, providing an overview of several approaches that aim to develop techniques for the (automatic or semi-automatic) verification of agent-based systems.

When using logical techniques, there are two common approaches to the verification of computational systems: model checking [CGP99] and theorem proving. Model checking essentially requires a finite-state description of the possible behaviours of the system, together with a property to be checked. A model checker then carries out an exhaustive search through the finite-state description attempting to find a way where the property can be falsified. If the model checker is unable to find such a counter-example, then the property is satisfied on *all* possible executions of the system. Model checking has become very popular over recent years, with theoretical advances in the use of model checking combined with significant efficiency improvements in model checking tools leading to a marked increase in the practical use of model checkers. Alternatively, theorem proving requires a logical formula that captures *all* the behaviours of the system (whether finite or not) and a property, expressed as a logical formula. Then a proof must be carried out, showing that the formula describing the property is a *logical consequence* of the formula describing the system. While model checking is usually automatic, theorem-proving usually involves human intervention. Between these two extremes are proof problems that can be solved automatically, yet involve infinite-state behaviours.

Typically, autonomous agents in a multi-agent system need to share and/or compete for resources. The usual approach is that resources are owned by specific agents and the allocation mechanism is based on some form of negotiation technique. Usually, the communication and allocation protocols used need themselves be verified as they could incur in well known faults (e.g., starvation). Indeed, any such protocols can be independently checked for properties that guarantee no such faults can occur. Ideally, we will have, in the near future, libraries of such high-level agent coordination mechanisms associated with properties that have already been verified. Developers of multi-agent systems could then use these “off the shelf”, while being aware of their properties (as well as the requirements with which agents should comply in order to ensure appropriate interaction).

Although still at an early stage, this points to the use of compositional reasoning for the formal verification of multi-agent systems. In addition, verification must take into account the use of *legacy* code, something which is often used in multi-agent systems. Further, it is important to note that verification of multi-agent systems can work on several levels of abstraction: the level of social interactions and social organisations; the individual level where the agent’s internal reasoning is targeted; and also an even lower level concerning specific algorithms or legacy code used internally to an agent.

While most of this research aims to reason on agent specifications or models with an a-priori approach, a different aspect of verification in multi-agent systems is run-time verification. This accounts to monitoring and checking whether a particular implemented, running agent does indeed operate according to its specification. There exist logic-based tools that reason upon the externally observable behaviour of interacting agents and verify whether it complies to predefined norms or protocols. In particular, the SOCS-SI tool [ACG⁺06] (already mentioned in the previous section in relation with SOCS), uses the abductive logic programming SCIFF proof-procedure to consider messages exchanged by agents plus other events and carry out such a run-time interaction verification task. Though comparatively less studied, this is also an important aspect of developing dependable agent systems.

However, while model checking techniques can be seen as appropriate solutions to many of the above problems, they are less satisfactory when we consider issues such as openness (i.e., an arbitrary number of autonomous agents may be involved) and emergence (i.e., coping with evolving agent behaviours).

Here, theorem-proving appears to have advantages over model-checking. Although it is early days, our view is that combining deductive [Fis05a] and algorithmic techniques (i.e., the various approaches based on model-checking discussed below) will have an important role in the verification of multi-agent systems in the future.

Regardless of how long it takes to achieve these more ambitious objective, in the short term we expect to see an increase in the number of researchers interested in formal verification of multi-agent systems, mainly because of the adequacy of such systems for many areas of application that demand large-scale dependable systems. In particular, we expect to see both theorem-proving and model-checking techniques, and hopefully combinations of both, being developed for the many logics of interest for multi-agent systems. Various examples are given below in this section, showing which are the logics of interest, and also the kinds of combined logics that are required, as most of those logics that appear in the literature so far cover only a part of the features that multi-agent systems normally have (such as epistemic properties, temporal properties, motivational properties, coalitional power, etc.).

4.1 Model Checking

Model checking essentially takes a finite state model of a program and a temporal property and then checks that the temporal property is satisfied on *all* possible executions of the program [CGP99]. In the case of agent verification, we again need a finite state agent description (often a program) and a property to be checked on all possible executions of the agent. In using model checking to verify properties of agent-based systems, we are often using properties based on some combination of modal and temporal logics, such as BDI logics, ATL, or temporal logics of knowledge. As in the standard case, problems of space requirements and of handling essentially infinite-state programs, are the key issues in agent verification. A further complication is that, once multi-agent systems are tackled, the complex interactions between such agents often mean the potential state space of the whole system is very large.

- **Model Checking Algorithms for BDI-like Logics**

In [BGS98], a model-checking algorithm for a BDI logic (using CTL as underlying temporal logic) was presented (see [RG98, Woo00, SRG99] for details on BDI-like logics). The approach formalises the BDI modalities using a variation of multi-language hierarchical logics [GS94, GG01]; the idea is to have a tree hierarchy of related theories, representing the nesting of modalities (that is, the mental model an agent has, possibly about another agent). In [BC02], the approach was extended to consider symbolic model-checking [CGP99] in particular (previous work had given algorithms for explicit-state model checking). One drawback of the approach is that it is not clear how the types of models they require are to be generated from actual systems, therefore their techniques might not be easily applicable to practical verification of multi-agent systems.

- **Model Checking AgentSpeak Programs**

In a series of papers, Bordini *et al.* have developed model-checking techniques that apply directly to multi-agent programs written in AgentSpeak. The approach is to translate AgentSpeak multi-agent systems into either Promela [BFP⁺03] or Java [BFV⁺04c] models, then using, respectively, SPIN [Hol03] or JPF [VHB⁺00] as model checkers. The property specification language is a much simplified BDI logic which is suitable for AgentSpeak programs. The tools that have been developed to automate the translation process were briefly described in [BVF⁺03]. In [BFV⁺04b], a slicing algorithm for AgentSpeak was introduced, and initial experimental results shown in that paper indicate that the technique can have a significant impact in reducing the state-space that needs to be considered in practical model-checking for AgentSpeak programs. The whole approach was summarised in [BFV⁺04a, BFV⁺06].

- **Bounded Model Checking for Epistemic Temporal Logics**

Penczek and Lomuscio have defined bounded semantics of CTLK [PL03], a combined logic of knowledge and time. Based on that semantics, they define a bounded model-checking algorithm for that logic. The approach is to translate the (bounded) model of the system (given as an interpreted system) and a formula ϕ — the (negation of the) property to be verified — to sets of propositional formulae in such a way that they can only be satisfied if, and only if, the system satisfies ϕ . This means that, in their approach, a model checking exercise can be carried out by existing SAT-solvers (which can at times, they argue, turn out to be rather efficient). More recently [RL04], they have defined similar translations to OBDDs, the approach used for symbolic model checking [CGP99].

Experimental results on the “attacking generals” problem, which has applications for agent coordination, were reported in [LLP04]. It remains to be seen if the performance of both SAT- and OBDD-based model-checking (and the associated translations) for practical verification of multi-agent systems will be satisfactory.

- **Model Checking for (variations of) ATL**

The emphasis in ATL verification is usually on model checking techniques even though, with complete ATL axiomatizations [GvD06, WLW⁺06], there is potential for deductive approaches. Another logic with a game-theoretic basis that has been used for reasoning about teams, coalitions, etc., is Coalition Logic [Pau02], which can be seen as a generalisation of the modal base of Game Logic [Par83, PR03] (itself an extension of Propositional Dynamic Logic [Har00]). Practical model checking techniques for ATEL [vdHW03], a version of ATL with added epistemic modalities, were presented in [vOvdHW04] (by reduction to ATL); the paper includes an exercise in model checking the Russian cards scenario using Mocha [AdAH⁺00], an ATL model checker.

Further work based on ATL for agent verification has been carried out by Ryan and Schobbens [RS02]. They define a refinement relation between agents in a multi-agent system based on ATL. This allows the assume-guarantee paradigm (a well-known form of compositional reasoning [CGP99]) to be used in combination with model checking, thus extending (in principle) the use of formal verification methods to much larger systems, despite the state-space explosion problem.

- **Deontic Interpreted Systems**

Interpreted systems have arguably a more clear correspondence with states of distributed computational systems than Kripke structures, the preferred semantic formalism in philosophical logic. Lomuscio and Sergot [LS02] have extended interpreted systems with deontic notions [Aqv83]; in particular, this allows one to reason about correct and incorrect behaviour of individual agents and of the multi-agent system. In [LS01], they extended the framework so as to combine it with standard epistemic notions [FHM⁺95], thus allowing the expression of properties such as the knowledge that agents have, or require, when they function properly (with regards to a certain protocol). See also [LS03] for an extended paper on the subject. Regarding actual verification, Penczek and Lomuscio have applied bounded model-checking techniques to deontic interpreted systems [WPL05].

- **Checking Communication Protocols within Electronic Institutions**

Cliffe and Padget [CP02] introduced a language that is suitable for modelling communication protocols in the context of Electronic Institutions [ERS⁺01]. The language used to write the models of communication protocols provides a Prolog-like notation for defining a transition system, which is then translated into the input language of NuSMV [CCG⁺02], a well-known symbolic model checker (model checking is then carried out by NuSMV). It is not yet clear how applicable the approach will be in general.

- **Model Checking Knowledge-Based Programs**

In [vdM98], van der Meyden examined the model checking problem for epistemic logic (with common knowledge formulæ) and systems of multiple agents (with perfect recall) situated in finite state environments. Model checking techniques for a logic combining knowledge and temporal modalities are described in [vdMW03]. This is the underlying theory for the MCK model checker [GvdM04], available at <http://www.cse.unsw.edu.au/~mck/>. MCK is a symbolic model checker implemented in Haskell, but relying on David Long's BDD package which is implemented in C (<http://www.cs.cmu.edu/~modelcheck/bdd.html>).

- **Model Checking Agent Dialogues**

A language for describing agent dialogues, called MAP and based on CCS, was defined by Walton in [Wal05]. The language is particularly targeted at modelling dialogues in the context of Electronic Institutions [ERS⁺01] (the same application area as work by Cliffe & Padget described earlier). As in the work on model checking for AgentSpeak described above, the approach here is to translate MAP protocols into Promela and use SPIN as model checker. Another use of the MAP language, for verifying agent communication in interaction with web services, is mentioned in [Wal04].

- **Verifying Commitments**

Both Singh and Wooldridge have argued about the difficulty of checking compliance of agents to communication protocols because of the mentalistic approach used in giving semantics to agent communication languages [Sin98, Woo98]. An alternative to the mentalistic approach is to use social commitments as the basis for speech-act semantics [Sin98]. In [YS02, YS04], Yolum and Singh present an approach to representing and reasoning about commitment protocols, based on the Event Calculus [KS86]. They use an event calculus planner to allow agents to choose flexibly a course of action that is compliant with given protocols. Another approach to commitment protocols is presented in [WS04]. The idea there is to map Dooley graphs (often used for designing multi-agent interaction protocols) and commitment causal relations to the π -calculus [MPW92], which can then be used to check consistency of multi-agent systems.

4.2 Theorem Proving

In contrast to model checking, the deductive approach to agent verification uses a logical formula to describe all possible executions of the agent system, and then attempts to *prove* the required property from this logical formula. Again, in considering rational agents, the properties required are often captured using combinations of modal and temporal logics, therefore the deduction problem can be quite difficult. As usual with theorem proving, there are two main varieties: those that involve decidable problems, where the deduction is guaranteed to terminate, and much more complex problems, where human interaction is required.

- **Temporal Logics**

In situations where agent descriptions are solely given using temporal logics (e.g., METATEM above), then temporal proof may be carried out. In general, first order temporal problems are very complex, but there are automatic methods for both propositional [FDP01, HK03] and significant fragments of first-order [DFK06, KDDFH05, HKRV04] temporal logics. While individual modal, temporal, or description logics have been studied extensively in the past, much less is known about combinations of modal logics and, in particular, about practical proof methods for such combinations [BDF⁺01]. However, both resolution- and tableau-based proof methods

for combinations of propositional (linear or branching time) temporal logics with modal logics [DFW98, DFB02, WDF98, HDSF00] have been studied. For example, the combination of modal and temporal logics and translations into fragments of classical first-order logics have been used to reason about the KARO agent framework [HDS⁺06], and about extensions of the basic METATEM framework [Fis05a].

- **CASL**

Shapiro, Lespérance, and Levesque [SLL02] introduced the Cognitive Agent Specification Language (CASL), which is based on the situation calculus, knowledge and goal operators (for representing an agent's mental states), and uses ConGolog (see above) for describing an agent's behaviour. They also defined CASL_{ve}, a verification environment for CASL that consists of representing a CASL specification within the framework of the PVS verification system [ORS92].

- **BDI and Logics of Concurrency** Schild [Sch00] provides a mapping from Rao and Georgeff's BDI logic [RG98] to μ -calculus [Koz83]. The BDI model of Rao and Georgeff is based on a combination of the branching time logic CTL* [ES89] and modal operators for belief, desire, and intention. While it had been shown that CTL* is subsumed by the μ -calculus [Eme90], it was much less clear how to deal with the orthogonality that Rao and Georgeff's logic exhibits. Schild tackles this by collapsing the two dimensions of time and modalities onto a one dimensional structure. The translation into μ -calculus provides a full axiomatisation of Rao and Georgeff's BDI logic.

Schild proves that, for CTLⁿ (containing formulae with at most n nested branching operators), the translation is decidable and succinct. Some axiom schemata (such as weak realism) can also be decided, while others (such as strong realism) might yield non-succinct translations, and are therefore less usable for automatic verification.

5 Concluding Remarks

5.1 Related Work

While this paper focusses on the use of computational logic within research on agents, there are many non-logical approaches to programming agents. Without claiming to cover all (important) frameworks, we briefly review some of the most prominent ones.

- **FIPA-OS** [<http://fipa-os.sourceforge.net>]

FIPA-OS is a component-based toolkit enabling rapid development of FIPA compliant agents. FIPA-OS supports the majority of the FIPA experimental specifications and is being continuously improved as a managed *Open Source* community project, making it a suitable choice on which to base FIPA-compliant agent development activity. FIPA-OS was first released in August 1999 as the first publicly available implementation of FIPA technology.

- **PlanGent** [<http://www.toshiba.co.jp/plangent>]

PlanGent is an intelligent agent system that performs tasks for human users. An agent in the PlanGent system has movement and planning capabilities; it can move around the network, determine its best course of action in various situations, and act autonomously.

- **JACK** [<http://www.agent-software.com>]

JACK Intelligent Agents is an environment for building, running, and integrating commercial-grade multi-agent systems using a component-based approach. This approach is based heavily on

the BDI architecture and appropriate Java classes for agent programming are provided. The JACK Agent Language is the programming language that extends Java with agent-oriented concepts, such as agents, capabilities, events, plans, agent knowledge bases (databases), and resource and concurrency management.

- **Jade** [<http://jade.tilab.com/>]

JADE (Java Agent DEvelopment Framework) is a software framework implemented in Java. It supports FIPA-based communication and provides a set of graphical tools that aid the programmer during multi-agent systems development. In recent years, JADE has become one of the most used agent platforms within the research community, and its middleware is quickly becoming a *de facto* standard. An extension to JADE called Jadex extends the basic framework with a BDI-like architecture (<http://vsis-www.informatik.uni-hamburg.de/projects/jadex/>).

- **Aglets** [<http://aglets.sourceforge.net/>]

Aglets have been developed by the IBM Tokyo Research Laboratory, and provide an applet-like approach to building agents using Java. It consists of a daemon application which serves as a basic framework, and mobile agents which can communicate using the framework. Through the use of Java security mechanisms, aglets support advanced security features.

- **Cougaar** [<http://www.cougaar.org/>]

Cougaar is a Java-based framework for the creation of large-scale distributed applications. Originally conceived as part of a DARPA project, it has since become an open-source project. Special concern has been “survivability” in case of failures in parts of the system, as well as security. Based on a flexible component model, it implements various communication approaches, such as blackboard, and various protocols, such as CORBA, RMI, HTTP, and UDP.

5.2 Frequently Asked Questions

Q: Why should we use logical approaches to agents?

A: Two reasons. The first is clarity – logical descriptions are unambiguous and often concise, thus providing a basis for detailed analysis and verification. The second is that, when using logical approaches, there is a vast number of different logics (or combinations) to choose from. This gives the system designer a very wide range of possible approaches. Choosing the best logic can give very intuitive and useful abstractions for programming complex systems.

Q: Have the “right” logics been used?

A: Unfortunately, there is no “right” logic! In most cases, particular logical mechanisms have been developed for the application area being considered. Even the BDI approach, which is probably the most widespread in this area, has been modified, or even avoided, for certain applications.

Q: Logics are complex, so isn’t using them likely to be inefficient?

A: Yes, and no! Logical manipulation is, in general, complex. But, in many cases, specific and efficient logical fragments have been devised.

Q: What about industry?

A: It is true that logical approaches to multi-agent systems are not widely used in industry (yet). However, as concerns increase over the reliability and security of complex systems and over the public's *trust* in these systems, so the use of logical approaches are likely to increase.

Q: Is there a killer application?

A: No, but see below.

5.3 Future Directions

Below is a subjective and, to some extent, speculative view about the future of computational logics and agents.

Agent Programming We have yet to see a truly successful agent programming language. Often agents are still programmed just using Java. Thus, there is an opportunity for a computational logic approach to programming agents to take a lead here.

Agent Analysis Agent verification will surely increase. When unknown agents are sent to you (e.g., in the form of Java bytecodes) and your agents are expected to negotiate (sometimes sensitive) information with unknown agents, then verification will be increasingly required for questions of *security* and *trust*.

Agent Applications We conclude the paper by suggesting four areas where computational logic approaches to agency may be particularly influential.

1. Ubiquitous Computing

With the mobility, openness, and organisational aspects, both programming and analysing ubiquitous computing applications can be problematic. By providing a logical basis for programming both individual agents (through key aspects such as deliberation) and multi-agent systems (through groups and teams), computational logic can help the principled development of reliable ubiquitous computing applications. In addition, logical verification techniques can be used to analyse the behaviour of such potentially complex applications.

2. Autonomous Control

Control systems are highly developed. However, the requirement for systems to act autonomously and *rationally* does not necessarily sit well with traditional control systems. As we move from systems where we need to know *what* the system does, to those where we need to know *why* it did something, we believe that logical approaches to the control and analysis of complex systems will come into their own.

3. Semantic Web

While we have said very little about the Semantic Web in this report, web services are complementary to agent technology. Consequently, many techniques for agent-based systems can be transferred to the Semantic Web. In terms of computational logic, this is even more evident when we see that description logics are at the core of Web ontologies. Thus, computational logic has a key role to play both in programming web-service description and discovery, and in analysing web-service collections.

4. Distribution

This covers a wide variety of aspects, from negotiation, through resource allocation/balancing and stock markets, to e-commerce and financial applications. In many of these cases, varieties of economic games are used. However, there is growing indication that computational logic might take a more central role in this area, possibly through combined logics/economics approaches.

References

- [AdAH⁺00] R. Alur, L. de Alfaro, T. A. Henzinger, S. C. Krishnan, F. Y. C. Mang, S. Qadeer, S. K. Rajamani, and S. Taşiran. *MOCHA user manual*. University of Berkeley Report, 2000.
- [ABH⁺06] N. Alechina, R. Bordini, J. Hübner, M. Jago, and B. Logan. Automating Belief Revision for AgentSpeak. In M. Baldoni and U. Endriss, editors, *Declarative Agent Languages and Technologies IV*, LNAI, Heidelberg, 2006. Springer-Verlag.
- [ACG⁺04a] M. Alberti, F. Chesani, M. Gavanelli, E. Lamma, P. Mello, and P. Torroni. A Logic Based Approach to Interaction Design in Open Multi-agent Systems. In *Proc. Thirteenth IEEE international Workshops on Enabling Technologies: Infrastructures for Collaborative Enterprises (WETICE), Second International Workshop "Theory and Practice of Open Computational Systems (TAPOCS)"*, pages 387–392, Modena, Italy, June 14–16 2004. IEEE Press.
- [ACG⁺04b] M. Alberti, F. Chesani, M. Gavanelli, E. Lamma, P. Mello, and P. Torroni. The SOCS computational logic approach for the specification and verification of agent societies. In *Global Computing Workshop*, volume 3267 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, 2004.
- [ACG⁺06] M. Alberti, F. Chesani, M. Gavanelli, E. Lamma, P. Mello, and P. Torroni. Compliance Verification of Agent Interaction: A Logic-based Software Tool. *Applied Artificial Intelligence*, 20(2-4):133-157, Taylor & Francis, April 2006.
- [ADG⁺04] M. Alberti, D. Daolio, M. Gavanelli, E. Lamma, P. Mello, and P. Torroni. Specification and verification of agent interaction protocols in a logic-based system. In *Proc. Nineteenth Annual ACM Symposium on Applied Computing (SAC). Special Track on Agents, Interactions, Mobility, and Systems (AIMS)*, pages 72–78, Nicosia, Cyprus, March 14–17 2004. ACM Press.
- [AGL⁺03a] M. Alberti, M. Gavanelli, E. Lamma, P. Mello, and P. Torroni. An Abductive Interpretation for Open Societies. In *Advances in Artificial Intelligence, Proc. Eighth Congress of the Italian Association for Artificial Intelligence (AI*IA)*, volume 2829 of *Lecture Notes in Artificial Intelligence*, pages 287–299. Springer-Verlag, September 23–26 2003.
- [AGL⁺03b] M. Alberti, M. Gavanelli, E. Lamma, P. Mello, and P. Torroni. Specification and verification of agent interactions using social integrity constraints. *Electronic Notes in Theoretical Computer Science*, 85(2), 2003.
- [AGL⁺04] M. Alberti, M. Gavanelli, E. Lamma, P. Mello, and P. Torroni. Modeling interactions using *Social Integrity Constraints*: A resource sharing case study. In *Declarative Agent Languages and Technologies*, volume 2990 of *Lecture Notes in Artificial Intelligence*, pages 243–262. Springer-Verlag, May 2004.
- [AGL⁺05] M. Alberti, M. Gavanelli, E. Lamma, P. Mello, and P. Torroni. The SCIFF abductive proof-procedure. In *Advances in Artificial Intelligence. Proc. Ninth Italian Congress on Artificial Intelligence (AI*IA)*, volume 3673 of *Lecture Notes in Artificial Intelligence*, pages 135–147. Springer-Verlag, 2005.
- [AGM85] C. E. Alchourron, P. Gardenfors and D. Makinson. On the Logic of Theory Change: Partial Meet Functions for Contraction and Revision. *Journal of Symbolic Logic*, 50:510–530, 1985.
- [AHK97] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-Time Temporal Logic. In *38th Annual Symposium on Foundations of Computer Science (FOCS '97), Miami Beach, Florida, USA, October 19-22*, pages 100–109. IEEE Computer Society, 1997.

- [AHK02] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-Time Temporal Logic. *Journal of the ACM*, 49(5):672–713, 2002.
- [AJL06] N. Alechina, M. Jago, and B. Logan. Resource-bounded Belief Revision and Contraction. In M. Baldoni, U. Endriss, A. Omicini, and P. Torroni, editors, *Declarative Agent Languages and Technologies III*, volume 3904 of *LNAI*, pages 141–154, Heidelberg, 2006. Springer-Verlag.
- [AL04] J. J. Alferes and J. A. Leite, editors. *Proc. European Conference on Logics in Artificial Intelligence (JELIA)*, volume 3229 of *Lecture Notes in Computer Science*. Springer, 2004.
- [ALF] ALFEBIITE: A Logical Framework for Ethical Behaviour between Infohabitants in the Information Trading Economy of the Universal Information Ecosystem. IST-1999-10298.
Home Page: <http://www.iis.ee.ic.ac.uk/~alfebiite/ab-home.htm>.
- [ALI] ALIAS: Abductive Logic Agent System.
Home Page: <http://lia.deis.unibo.it/Research/ALIAS/>.
- [ALP⁺98] J. J. Alferes, J. A. Leite, L. M. Pereira, H. Przymusinska, and T. C. Przymusinski. Dynamic logic programming. In *Proc. Sixth International Conference on Principles of Knowledge Representation and Reasoning*. Morgan Kaufmann Publishers, 1998.
- [AOR⁺99] K. A. Arisha, F. Ozcan, R. Ross, V. S. Subrahmanian, T. Eiter, and S. Kraus. IMPACT: a Platform for Collaborating Agents. *IEEE Intelligent Systems*, 14(2):64–72, March/April 1999.
- [APP⁺99] J. J. Alferes, L. M. Pereira, H. Przymusinska, and T. C. Przymusinski. LUPS: a language for updating logic programs. In *Proc. Fifth International Conference on Logic Programming and Non-Monotonic Reasoning (LPNMR)*. Volume 1730 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, 1999.
- [APS02] A. Artikis, J. Pitt, and M. Sergot. Animated Specifications of Computational Societies. In *Proc. First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1053–1061, Bologna, Italy. ACM Press, 2002.
- [Aqv83] L. Aqvist. Deontic Logic. In Gabbay and Guenther (eds), *Handbook of Philosophical Logic, Vol. II* D Reidel Publishing Company, pp.605-714, 1983.
- [Bak91] A. B. Baker. Nonmonotonic Reasoning in the Framework of Situation Calculus. *Artificial Intelligence*, 49(1–3):5–23, 1991.
- [Bar03] C. Baral. *Knowledge Representation, Reasoning and Declarative Problem Solving with Answer Sets*. Cambridge University Press, 2003.
- [BC02] M. Benerecetti and A. Cimatti. Symbolic Model Checking for Multi-Agent Systems. In *Proc. Workshop on Model Checking and Artificial Intelligence (MoChArt), Lyon, France, 22–23 July, 2002*.
- [BCM^c+03] F. Baader, D. Calvanese, D. M^cGuinness, D. Nardi, and P. Patel-Schneider, editors. *The Description Logic Handbook*. Cambridge University Press, 2003.
- [BDD⁺05] R. H. Bordini, M. Dastani, J. Dix, and A. El Fallah Seghrouchni, editors. *Multi-Agent Programming: Languages, Platforms and Applications*. Springer-Verlag, 2005.
- [BDE⁺04] A. Bracciali, N. Demetriou, U. Endriss, A. Kakas, W. Lu, P. Mancarella, F. Sadri, K. Stathis, F. Toni, and G. Terreni. The KGP Model of Agency: Computational Model and Prototype Implementation. In *Global Computing Workshop*, volume 3267 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, 2004.
- [BDF⁺01] B. Bennett, C. Dixon, M. Fisher, E. Franconi, I. Horrocks, and M. de Rijke. Combinations of Modal Logics. *AI Review*, 17(1):1–20, 2002.
- [BDH⁺01] J. Broersen, M. Dastani, Z. Huang, J. Hustijn, and L. van der Torre. The BOID Architecture. In *Proc. Fifth International Conference on Autonomous Agents*, pages 9–16. ACM Press, 2001.
- [BDH⁺02] J. Broersen, M. Dastani, Z. Huang, J. Hustijn, and L. van der Torre. Goal Generation in the BOID Architecture. In *Cognitive Science Quarterly*, volume 2, pages 428–447. 2002.

- [BED⁺06] A. Bracciali, U. Endriss, N. Demetriou, A. C. Kakas, W. Lu, and K. Stathis. Crafting the Mind of PROSOCS Agents. *Applied Artificial Intelligence*, 20(2-4):105-131. Taylor & Francis, 2006.
- [BFG⁺95] H. Barringer, M. Fisher, D. Gabbay, G. Gough, and R. Owens. METATEM: An Introduction. *Formal Aspects of Computing* 7(5):533–549. Springer-Verlag, 1995.
- [BFG⁺96] H. Barringer, M. Fisher, D. Gabbay, R. Owens, and M. Reynolds. *The Imperative Future: Principles of Executable Temporal Logics*. Research Studies Press, Chichester, United Kingdom, 1996.
- [BFP⁺03] R. H. Bordini, M. Fisher, C. Pardavila, and M. Wooldridge. Model checking AgentSpeak. In *Proc. Second International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 409–416. ACM Press, 2003.
- [BFV⁺04a] R. H. Bordini, M. Fisher, W. Visser, and M. Wooldridge. Model Checking Rational Agents. *IEEE Intelligent Systems*, 19(5):46–52, September/October 2004.
- [BFV⁺04b] R. H. Bordini, M. Fisher, W. Visser, and M. Wooldridge. State-Space Reduction Techniques in Agent Verification. In *Proc. Third International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 896–903. ACM Press, 2004.
- [BFV⁺04c] R. H. Bordini, M. Fisher, W. Visser, and M. Wooldridge. Verifiable Multi-Agent Programs. In *Proc. First International Workshop on Programming Multi-Agent Systems (ProMAS)*. Volume 3067 of *Lecture Notes in Artificial Intelligence*, pages 72–89. Springer-Verlag, 2004.
- [BFV⁺06] R. H. Bordini, M. Fisher, W. Visser, and M. Wooldridge. Verifying Multi-Agent Programs by Model Checking. *Journal of Autonomous Agents and Multi-Agent Systems*, 12(2):239–256, March 2006.
- [BGS98] M. Benerecetti, F. Giunchiglia, and L. Serafini. Model Checking Multiagent Systems. *Journal of Logic and Computation*, 8:401–423, 1998.
- [BHV05] R. H. Bordini, J. F. Hübner, and R. Vieira. **Jason** and the golden fleece of agent-oriented programming. In Bordini et al. [BDD⁺05], chapter 1, pages 3–37.
- [BMST05] A. Bracciali, P. Mancarella, K. Stathis, and F. Toni. On modelling declaratively multi-agent systems. In *Declarative Agent Languages and Technologies (DALT)*, volume 3476 of *Lecture Notes in Artificial Intelligence*, pages 53–68. Springer-Verlag, 2005.
- [Bor06] R. H. Bordini. BDI agent programming in AgentSpeak using **jason** (tutorial paper). In Toni and Torroni [TT06], pages 143–164.
- [Bra87] M. E. Bratman. *Intentions, Plans, and Practical Reason*. Harvard University Press, Cambridge, MA, 1987.
- [BVF⁺03] R. H. Bordini, W. Visser, M. Fisher, C. Pardavila, and M. Wooldridge. Model Checking Multi-Agent Programs with CASP. In *Proc. Fifteenth International Conference on Computer-Aided Verification (CAV)*. Volume 2725 of *Lecture Notes in Computer Science*, pages 110–113. Springer-Verlag, 2003. [Tool description]
- [CCG⁺02] A. Cimatti, E. M. Clarke, E. Giunchiglia, F. Giunchiglia, M. Pistore, M. Roveri, R. Sebastiani, and A. Tacchella. NuSMV2: An OpenSource Tool for Symbolic Model Checking. In *Proc. Fourteenth International Conference on Computer Aided Verification (CAV)*, volume 2404 of *Lecture Notes in Computer Science*, pages 359–364, Berlin, 2002. Springer Verlag.
- [CGA⁺06] F. Chesani, M. Gavanelli, M. Alberti, E. Lamma, P. Mello, and P. Torroni. Specification and verification of agent interaction using abductive reasoning (tutorial paper). In Toni and Torroni [TT06], pages 243–264.
- [CGP99] E. M. Clarke, Jr., O. Grumberg, and D. A. Peled. *Model Checking*. The MIT Press, 1999.
- [CLM⁺03] A. Ciampolini, E. Lamma, P. Mello, F. Toni, and P. Torroni. Co-operation and Competition in **ALIAS**: A Logic Framework for Agents that Negotiate. *Annals of Mathematics and Artificial Intelligence*, 37(1-2):65–91, 2003.
- [CLMT02] A. Ciampolini, E. Lamma, P. Mello, and P. Torroni. **LAILA**: A Language for Coordinating Abductive Reasoning among Logic Agents. *Computer Languages*, 27(4):137–161, February 2002.

- [CP02] O. Cliffe and J. Padget. A Framework for Checking Agent Transactions within Institutions. In *Proc. Workshop on Model Checking and Artificial Intelligence (MoChArt-2002)*, Lyon, France, 2002.
- [CRH01] K. L. Clark, P. Robinson, and R. Hagen. Multi-threading and Message Communication in Qu-Prolog. *Theory and Practice of Logic Programming*, 1(3):283–301, 2001.
- [CT02] S. Costantini and A. Tocchio. A Logic Programming Language for Multi-Agent Systems. In *Proc. Eighth European Conference on Logics in Artificial Intelligence (JELIA)*, volume 2424 of *Lecture Notes in Computer Science*, pages 1–13. Springer-Verlag, September 2002.
- [CT04] S. Costantini and A. Tocchio. The DALI Logic Programming Agent-Oriented Language. In Alferes and Leite [AL04], pages 685–688.
- [CRZ06] K.L. Clark, P.J. Robinson, and S. Zappacosta-Amboldi. Multi-threaded Communicating Agents in Qu-Prolog (tutorial paper). In Toni and Torroni [TT06], pages 186–205.
- [DAML-S] A. Ankolenkar, M. Burstein, J. R. Hobbs, O. Lassila, D. L. Martin, D. McDermott, S. A. McIlraith, S. Narayanan, M. Paolucci, T. R. Payne, and K. Sycara. DAML-S: Web Service Description for the Semantic Web. In *Proc. First International Semantic Web Conference (ISWC)*, 2002.
- [DdBDM03] M. Dastani, F. de Boer, F. Dignum and J.-J. Ch Meyer. Programming Agent Deliberation: An approach illustrated using the 3APL language. In *Proc. Second International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 2003.
- [DFB02] C. Dixon, M. Fisher, and A. Bolotov. Clausal Resolution in a Logic of Rational Agency *Artificial Intelligence*, 139(1)47-89. Elsevier Science, 2002.
- [DFK06] A. Degtyarev, M. Fisher, and B. Konev, Monodic Temporal Resolution. *ACM Transactions on Computational Logic*, 7(1):108–150, January 2006.
- [DFW98] C. Dixon, M. Fisher, and M. Wooldridge. Resolution for Temporal Logics of Knowledge. *Journal of Logic and Computation*, 8(3):345–372, 1998.
- [Dig99] F. Dignum. Autonomous Agents and Norms. In *Artificial Intelligence and Law 7*:69–79, 1999.
- [dKLW98] M. d’Inverno, D. Kinny, M. Luck, and M. Wooldridge. A Formal Specification of dMARS. In Singh, Rao and Wooldridge [SRW98].
- [dL98] M. d’Inverno and M. Luck. Engineering AgentSpeak(L): A Formal Computational Model. In *Journal of Logic and Computation*, 8(3):233–260, 1998.
- [DK95] Y. Dimopoulos and A. C. Kakas. Logic Programming without Negation as Failure. In *Proc. International Symposium on Logic Programming*, pages 369–384, 1995.
- [DNP02] P. Dell’Acqua, U. Nilsson, and L. M. Pereira. A Logic Based Asynchronous Multi-Agent System. In *Electronic Notes in Theoretical Computer Science*, volume 70. Elsevier, 2002.
- [DP99] P. Dell’Acqua and L. M. Pereira. Updating Agents. In *Proc. ICLP Workshop on Multi-Agent Systems in Logic Programming (MAS)*, Las Cruces, New Mexico, 1999.
- [DP03] P. Dell’Acqua and L. M. Pereira. A Logical Framework for Modelling eMAS. In *Proc. Fifth International Symposium on Practical Aspects of Declarative Languages (PADL)*, volume 2562 of *Lecture Notes in Computer Science*, pages 241–255. Springer-Verlag, January 2003.
- [Eme90] E. A. Emerson. Temporal and Modal Logic. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B, chapter 16, pages 997–1072. Elsevier Science, Amsterdam, 1990.
- [EMS02] T. Eiter, V. Mascardi, and V. S. Subrahmanian. Error-Tolerant Agents. In *Computational Logic. Logic Programming and Beyond*, volume 2407 of *Lecture Notes in Computer Science*, pages 586–625. Springer-Verlag, 2002.
- [EMS⁺04] U. Endriss, P. Mancarella, F. Sadri, G. Terreni, and F. Toni. The CIFF Proof Procedure for Abductive Logic Programming with Constraints. In Alferes and Leite [AL04], pages 31–43.
- [ERS⁺01] M. Esteva, J. A. Rodríguez, C. Sierra, P. Garcia, and J. L. Arcos. On the formal specification of electronic institutions. In *Agent-mediated Electronic Commerce (The European AgentLink Perspective)*, volume 1991 of *Lecture Notes in Artificial Intelligence*, pages 126–147, 2001.

- [ES89] E. A. Emerson and J. Srinivasan. Branching Time Temporal Logic. In *Linear Time, Branching Time, and Partial Order in Logics and Models for Concurrency*, pages 123–172. Springer-Verlag, Berlin, 1989.
- [ES98] T. Eiter and V. S. Subrahmanian. Deontic Action Programs. In *Workshop on Foundations of Models and Languages for Data and Objects*, pages 37–54, 1998.
- [ES99] T. Eiter and V. S. Subrahmanian. Heterogeneous Active Agents, II: Algorithms and Complexity. *Artificial Intelligence*, 108(1–2):257–307, 1999.
- [ESP99] T. Eiter, V.S. Subrahmanian, and G. Pick. Heterogeneous Active Agents, I: Semantics. *Artificial Intelligence*, 108(1-2):179–255, March 1999.
- [ESR00] T. Eiter, V. S. Subrahmanian, and T. J. Rodgers. Heterogeneous Active Agents, III: Polynomially Implementable Agents. *Artificial Intelligence*, 117(1):107–167, 2000.
- [FDP01] M. Fisher, C. Dixon, and M. Peim. Clausal Temporal Resolution. *ACM Transactions on Computational Logic*, 2(1):12–56, January 2001.
- [FG99] M. Fisher and C. Ghidini. Programming Resource-Bounded Deliberative Agents. In *Proc. International Joint Conference on Artificial Intelligence (IJCAI)*, pages 200–205, 1999.
- [FG02] M. Fisher and C. Ghidini. The ABC of Rational Agent Modelling. In *Proc. First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 849–856. ACM Press, 2002.
- [FGH03] M. Fisher, C. Ghidini, and B. Hirsch. Organising Logic-Based Agents. In *Proc. Second NASA/IEEE Goddard Workshop on Formal Approaches to Agent Based Systems (FAABS)*, volume 2699 of *Lecture Notes in Computer Science*, pages 15–27, Berlin, 2003. Springer-Verlag.
- [FGH04] M. Fisher, C. Ghidini, and B. Hirsch. Programming Groups of Rational Agents. In *Computational Logic in Multi-Agent Systems (CLIMA-IV)*. Volume 3259 of *Lecture Notes in Computer Science*, Springer-Verlag, November 2004.
- [FGV05] M. Fisher, D. Gabbay and L. Vila, editors. *Handbook of Temporal Reasoning in Artificial Intelligence*. Elsevier Science, 2005.
- [FHM⁺95] R. Fagin, J. Y. Halpern, Y. Moses, and M. Vardi. *Reasoning about Knowledge*. The MIT Press, Cambridge, MA, 1995.
- [Fis94] M. Fisher. A Survey of Concurrent METATEM — The Language and its Applications. In *Proc. First International Conference on Temporal Logics*, volume 827 of *Lecture Notes in Artificial Intelligence*, pages 480–505. Springer-Verlag, Berlin, 1994.
- [Fis95] M. Fisher. Representing and Executing Agent-Based Systems. In Wooldridge and Jennings [WJ95b].
- [Fis96] M. Fisher. A Temporal Semantics for Concurrent METATEM. *Journal of Symbolic Computation* 22(5/6), November/December 1996.
- [Fis97a] M. Fisher. Implementing BDI-like Systems by Direct Execution. In *Proc. International Joint Conference on Artificial Intelligence (IJCAI)*, Morgan-Kaufmann, 1997.
- [Fis97b] M. Fisher. A Normal Form for Temporal Logic and its Application in Theorem-Proving and Execution. *Journal of Logic and Computation*, 7(4):429–456, August 1997.
- [Fis05a] M. Fisher. Temporal Development Methods for Agent-Based Systems. *Journal of Autonomous Agents and Multi-Agent Systems*, 10(1):41-66, January 2005.
- [Fis05b] M. Fisher. METATEM: The Story so Far. In *Proc. Third International Workshop on Programming Multiagent Systems (ProMAS)*, volume 3862 of *Lecture Notes in Artificial Intelligence*, pages 3-22, Springer Verlag, 2005.
- [FK97] T. H. Fung and R. A. Kowalski. The IFF proof procedure for Abductive Logic Programming. *Journal of Logic Programming*, 33(2):151–165, November 1997.

- [FW97] M. Fisher, and M. Wooldridge. On the Formal Specification and Verification of Multi-Agent Systems. *International Journal of Cooperative Information Systems*, 6(1):37–65, January 1997.
- [FW05] M. Fisher, and M. Wooldridge. Temporal Reasoning in Agent-Based Systems In [FGV05], pages 469–496.
- [GvD06] V. Goranko and G. van Drimmelen. Complete Axiomatisation and Decidability of Alternating-Time Temporal Logic. *Theoretical Computer Science*, 253(1-3):93–117, 2006.
- [GG01] C. Ghidini and F. Giunchiglia. Local Models Semantics, or Contextual Reasoning=Locality+Compatibility. *Artificial Intelligence*, 127(2):221–259, 2001.
- [GL87] M. P. Georgeff and A. L. Lansky. Reactive Reasoning and Planning. In *Proc. Sixth National Conference on Artificial Intelligence (AAAI)*, pages 677–682, 1987.
- [GL88] M. Gelfond and V. Lifschitz. The stable model semantics for logic programming. In *Proc. Joint International Conference and Symposium on Logic Programming*, pages 1070–1080. MIT Press, 1988.
- [GLL00] G. De Giacomo, Y. Lespérance, and H. J. Levesque. Congolog, a Concurrent Programming Language based on the Situation Calculus. *Artificial Intelligence* 121:109–169, 2000.
- [GvdM04] P. Gammie and R. van der Meyden. MCK: Model Checking the Logic of Knowledge. In *Proc. Sixteenth International Conference on Computer Aided Verification (CAV), Lecture Notes in Computer Science*, pages 479–483. Springer-Verlag, 2004.
- [Go192] R. Goldblatt. *Logics of Time and Computation*, volume 7 of *CLSI Lecture Notes*. CLSI, Stanford, CA, 2nd edition, 1992.
- [GS94] F. Giunchiglia and L. Serafini. Multilanguage Hierarchical Logics or: How we can do without modal logics. *Artificial Intelligence*, 65(1):29–70, 1994.
- [Har84] D. Harel. Dynamic Logic. In D. Gabbay and F. Guenther, editors, *Handbook of Philosophical Logic*, volume II, chapter 10, pages 497–604. D. Reidel Publishing Company, 1984.
- [Har00] D. Harel, D. Kozen, and J. Tiuryn. *Dynamic Logic*. The MIT Press, Cambridge, MA, 2000.
- [HdBdHM98] K.V. Hindriks, F.S. de Boer, W. van der Hoek and J.-J.Ch Meyer. Formal Semantics for an Abstract Agent Programming Language. In Singh, Rao and Wooldridge [SRW98], pages 215–229.
- [HdBHM01] K.V. Hindriks, F.S. de Boer, W. van der Hoek and J.-J.Ch Meyer. Agent Programming with Declarative Goals. In *Intelligent Agents VII*, volume 1986 of *Lecture Notes in Artificial Intelligence*, pages 228–243. Springer-Verlag 2001.
- [HDS⁺01] U. Hustadt, C. Dixon, R. A. Schmidt, M. Fisher, J-J. Meyer, and W. van der Hoek. Reasoning about Agents in the KARO framework. In *Proc. Eighth International Symposium on Temporal Representation and Reasoning (TIME)*, pages 206–213. IEEE Press, 2001.
- [HDS⁺06] U. Hustadt, C. Dixon, R. Schmidt, M. Fisher, J-J. Meyer and W. van der Hoek. Verification within the KARO Agent Theory. In *Agent Technology from a Formal Perspective*. Springer-Verlag, 2006.
- [HDSF00] U. Hustadt, C. Dixon, R. Schmidt, and M. Fisher. Normal Forms and Proofs in Combined Modal and Temporal Logics. In *Proc. Third International Workshop on Frontiers of Combining Systems (FroCoS)*, volume 1794 of *Lecture Notes in Artificial Intelligence*, pages 73–87. Springer, 2000.
- [HK03] U. Hustadt and B.Konev TRP++ 2.0: A Temporal Resolution Prover. In *Proc. International Conference on Automated Deduction (CADE)* pages 274–278. Volume 2741 of *Lecture Notes in Artificial Intelligence*, Springer, 2003.
- [HFGB05] B. Hirsch, M. Fisher, C. Ghidini, and P. Busetta. Organising Software in Active Environments. In *Proc. Workshop on Computational Logic in Multi-Agent Systems (CLIMA-V)*, volume 3487 of *Lecture Notes in Computer Science*. Springer-Verlag, 2005.
- [HKRV04] U. Hustadt, B. Konev, A. Riazanov, and A. Voronkov TeMP: A Temporal Monodic Prover. In *Proc. Second International Joint Conference on Automated Reasoning (IJCAR)*, pages 326–330. Volume 3097 of *Lecture Notes in Artificial Intelligence*. Springer, 2004.

- [vdHvLM93a] W. van der Hoek, B. van Linder, and J-J. Meyer. A Logic of Capabilities. Rapport IR-330, Vrije Universiteit, Amsterdam, July 1993.
- [vdHvLM93b] W. van der Hoek, B. van Linder, and John.-Jules Ch. Meyer. Unravelling Nondeterminism: On Having the Ability to Choose. Technical Report RUU-CS-93-3, Department of Computer Science, Utrecht University, 1993.
- [vdHvLM97] W. van der Hoek, B. van Linder, and J-J. Ch. Meyer. An Integrated Modal Approach to Rational Agents. Technical Report UU-CS-1997-06, Department of Computer Science, Utrecht University, 1997.
- [vdHMvS99] W. van der Hoek, J-J. Meyer, and J.W. van Schagen. Formalizing Potential of Agents. *Proc. Seventh CSLI Workshop on Logic, Language, and Computation*, 1999.
- [Hol03] G. J. Holzmann. *The Spin Model Checker: Primer and Reference Manual*. Addison-Wesley, 2003.
- [HP-S02] I. Horrocks, P. F. Patel-Schneider, and F. van Harmelen. Reviewing the Design of DAML+OIL: An Ontology Language for the Semantic Web. In *Proc. Eighteenth National Conference on Artificial Intelligence (AAAI)*, 2002.
- [HP-S03] I. Horrocks and P. F. Patel-Schneider. Reducing OWL Entailment to Description Logic Satisfiability. In *Proc. International Semantic Web Conference (ISWC)*, 2003.
- [HRL⁺97] H. Levesque, R. Reiter, Y. Lespérance, F. Lin and R. Scherl. GOLOG: A Logic Programming Language for Dynamic Domains. *Journal of Logic Programming* 31(1-3):59–84, 1997.
- [HW65] C. Hartshorne and P. Weiss, editors. *Collected Papers of Charles Sanders Peirce, 1931–1958*, volume 2. Harvards University Press, 1965.
- [vdHW03] W. van der Hoek and M. Wooldridge. Cooperation, Knowledge, and Time: Alternating-Time Temporal Epistemic Logic and its Applications. *Studia Logica*, 75(1):125–157, 2003.
- [JTW01] C. M. Jonker, J. Treur, and W. C. A. Wijngaards. Dynamics Within an Organisation: Temporal Specification, Simulation and Evaluation. In: *Proc. Third International Conference on Cognitive Science (ICCS)*, pp. 546-550. USTC Press, Beijing, 2001.
- [KKT93] A. C. Kakas, R. A. Kowalski, and F. Toni. Abductive Logic Programming. *Journal of Logic and Computation*, 2(6):719–770, 1993.
- [KS86] R. A. Kowalski and M. J. Sergot. A Logic-Based Calculus of Events. *New Generation Computing*, 4(1):67–95, 1986.
- [vL96] B. van Linder. *Modal Logics for Rational Agents*. PhD thesis, Universiteit Utrecht, 1996.
- [vLvDHM98] B. van Linder, W. van der Hoek, and J-J. Ch. Meyer. Formalising Abilities and Opportunities of Agents. *Fundamentae Informaticae*, 34(1,2):53–101, 1998.
- [vLMvdH97] B. van Linder, J-J. Ch. Meyer, and W. van der Hoek. Formalizing Motivational Attitudes of Agents using the KARO Framework. Technical Report UU-CS-1997-03, Utrecht University, Department of Computer Science, 1997.
- [KDDFH05] B. Konev, A. Degtyarev, C. Dixon, M. Fisher, and U. Hustadt, Mechanizing First-Order Temporal Resolution. *Information and Computation* 199(1–2):55–86. Elsevier Science, 2005.
- [KM02] A. C. Kakas and P. Moraitis. Argumentative Agent Deliberation, Roles and Context. In *Proc. Third International Workshop on Computational Logics in Multi-Agent Systems (CLIMA)*. Electronic Notes on Theoretical Computer Science, volume 70. Elsevier Science Publishers, 2002.
- [KM03] A. C. Kakas and P. Moraitis. Argumentation Based Decision Making for Autonomous Agents. In *Proc. Second International Joint Conference on Autonomous Agents and Multiagent Systems (AA-MAS)*, pages 883–890. ACM Press, 2003.
- [KMS⁺04] A. C. Kakas, P. Mancarella, F. Sadri, K. Stathis, and F. Toni. The KGP Model of Agency. In *Proc. Sixteenth European Conference on Artificial Intelligence (ECAI)*. IOS Press, August 2004.

- [Koz83] D. Kozen Results on the Propositional μ -calculus. In *Theoretical Computer Science*, 27:333—354, 1983
- [Kri63] S. Kripke. Semantical Considerations on Modal Logic. *Acta Philosophica Fennica*, 16:83–94, 1963.
- [KS99] R. A. Kowalski and F. Sadri. From Logic Programming towards Multi-Agent Systems. *Annals of Mathematics and Artificial Intelligence*, 25(3/4):391–419, 1999.
- [KTD04] A. C. Kakas, P. Torroni, and N. Demetriou. Agent Planning, Negotiation, and Control of Operation. In *Proc. Sixteenth European Conference on Artificial Intelligence (ECAI)*. IOS Press, August 2004, pages 28–32.
- [Lak99] G. Lakemeyer. On Sensing and Off-line Interpreting in GOLOG. In H. Levesque and F. Pirri, editors, *Logical Foundations for Cognitive Agents*, pages 173–189. Springer-Verlag Berlin, 1999.
- [LAP01] J. A. Leite, J. J. Alferes, and L. M. Pereira. Multi-dimensional dynamic knowledge representation. In *Proc. Sixth International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR)*, volume 2173 of *Lecture Notes in Artificial Intelligence*, pages 365–378. Springer-Verlag, 2001.
- [LAP02a] J. A. Leite, J. J. Alferes, and L. M. Pereira. *MINERVA*: A Dynamic Logic Programming Agent Architecture. In Meyer and Tambe [MT02], pages 141–157.
- [LAP⁺02b] J. A. Leite, J. J. Alferes, L. M. Pereira, H. Przymusinska, and T. C. Przymusinski. A Language for Multi-Dimensional Updates. In *Proc. Third International Workshop on Computational Logics in Multi-Agent Systems (CLIMA)*. *Electronic Notes on Theoretical Computer Science*, volume 70. Elsevier Science Publishers, 2002.
- [Lei03] J. A. Leite. *Evolving Knowledge Bases*. IOS Press, 2003.
- [LF94] Y. Labrou and T. Finin. A Semantics Approach for KQML—A General Purpose Communication Language for Software Agents. In *Proc. Third International Conference on Information and Knowledge Management (CIKM)*. ACM Press, 1994.
- [Lin95] F. Lin. Embracing Causality in Specifying the Indirect Effects of Actions. In *Proc. Fourteenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1985–1991, San Francisco, 1995. Morgan Kaufmann.
- [Llo87] J. W. Lloyd. *Foundations of Logic Programming*. Springer-Verlag, 2nd extended edition, 1987.
- [LLP04] A. Lomuscio, T. Lasica, and W Penczek. Bounded Model Checking for Interpreted Systems: Preliminary Experimental Results. In *Proc. Second NASA Workshop on Formal Approaches to Agent-Based Systems (FAABS)*, volume 2699 of *Lecture Notes in Computer Science*. Springer-Verlag, 2004.
- [LPR98a] H. Levesque, F. Pirri, and R. Reiter. Foundations for the Situation Calculus. *Linköping Electronic Articles in Computer and Information Science*, 3(18), 1998.
- [LPR98b] H. Levesque, F. Pirri, and R. Reiter. Foundations for the Situation Calculus. *Electronic Transactions on Artificial Intelligence* 2:159–178, 1998.
- [LS01] A. Lomuscio and M. Sergot. Extending Interpreted Systems with some Deontic Concepts. In *Theoretical Aspects of Rationality and Knowledge (TARK)*. Morgan Kauffman, 2001.
- [LS02] A. Lomuscio and M. Sergot. On Multi-Agent Systems Specification via Deontic Logic.
- [LS03] A. Lomuscio and M. Sergot. Deontic Interpreted Systems. *Studia Logica*, 75. 2003.
- [McC95] F. G. McCabe and K. L. Clark. April — Agent Process Interaction Language. In Wooldridge and Jennings [WJ95b], pages 324–340.
- [McC04] F. G. McCabe and K. L. Clark. Go! - A Multi-Paradigm Programming Language for Implementing Multi-Threaded Agents. *Annals of Mathematics in Artificial Intelligence*, 41(2-4):171-206, 2004.
- [McC86] J. McCarthy. Applications of circumscription to formalizing common-sense knowledge. *Artificial Intelligence*, 28:86–116, 1986.

- [McCH69] J. McCarthy and P. J. Hayes. Some Philosophical Problems from the Standpoint of Artificial Intelligence. In B. Meltzer and D. Michie, editors, *Machine Intelligence 4*, pages 463–502. Edinburgh University Press, 1969.
- [McGvH04] D. L. McGuinness and K. van Harmelen, editors. OWL Web Ontology Language — Overview. W3C Recommendation, 2004. Latest version available at <http://www.w3.org/TR/owl-features/>
- [MPW92] R. Milner, J. Parrow, and D. Walker. A calculus for mobile processes (parts I and II). *Information and Computation*, 100(1):1–40 and 41–77, September 1992.
- [MS01] S. McIlraith and T. Son. Adapting Golog for Programming in the Semantic Web. In *Proc. Fifth International Symposium on Logical Formalizations of Commonsense Reasoning*, pp 195–202. 2001.
- [MS02] S. McIlraith and T. Son. Adapting Golog for Composition of Semantic Web Services. In *Proc. International Conference on Principles of Knowledge Representation and Reasoning (KR&R)*, 2002.
- [vdM98] R. van der Meyden. Common Knowledge and Update in Finite Environments. *Information and Computation*, 140(2):115–157, Feb 1998.
- [MdBvE⁺01] J.-J. Ch. Meyer, F. de Boer, R. van Eijk, K. Hindriks, and W. van der Hoek. On Programming KARO Agents. *Logic Journal of the IGPL*, 9(22):261–272, 2001.
- [MD00] B. Mahony and J.S. Dong. Timed Communicating Object Z. *IEEE Transactions on Software Engineering* 26(2):150-177, Feb 2000.
- [MvdHvL99] J.-J. Meyer, W. van der Hoek, and B. van Linder. A Logical Approach to the Dynamics of Commitments. *Artificial Intelligence*, 113:1–40, 1999.
- [Moo84] R. C. Moore. A Formal Theory of Knowledge and Action. Technical Note 320, SRI International, February 1984.
- [Moo85] R. C. Moore. Semantical Considerations on Nonmonotonic Logic. *Artificial Intelligence*, 25:75–94, 1985.
- [Mor99] L. Morgenstern. Nonmonotonic logics. *MIT Encyclopedia of Cognitive Science*, 1999.
- [MT02] J.-J. Ch. Meyer and M. Tambe, editors, *Intelligent Agents VIII*, volume 2333 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, 2002.
- [vdMW03] R. van der Meyden and K. Wong. Complete Axiomatizations for Reasoning about Knowledge and Branching Time. *Studia Logica*, 75(1):93–123, Oct 2003.
- [MZW03] S. Merz, J. Zappe and M. Wirsing. A Spatio-Temporal Logic for the Specification and Refinement of Mobile Systems. In *Fundamental Approaches to Software Engineering (FASE)*, pp87–101, volume 2621 of *Lecture Notes in Computer Science*. Springer-Verlag, 2003.
- [vOvdHW04] S. van Otterloo, W. van der Hoek, and M. Wooldridge. Knowledge as strategic ability. In *Proc. First Workshop on Logic and Communication in Multi-Agent Systems (LCMAS)*. *Electronic Notes in Theoretical Computer Science*, 85(2), 2004.
- [ORS92] S. Owre, J. M. Rushby, and N. Shankar. PVS: A Prototype Verification System. In *Proc. Eleventh International Conference on Automated Deduction (CADE)*, volume 607 of *Lecture Notes in Artificial Intelligence* pages 748–752. Springer Verlag, 1992.
- [OWL] S. Bechhofer, F. van Harmelen, J. Hendler, I. Horrocks, D. L. McGuinness, P. F. Patel-Schneider, and L. A. Stein, editors. OWL Web Ontology Language Reference W3C Recommendation, 2004, <http://www.w3.org/TR/owl-ref/>
- [Par83] R. Parikh. Propositional Game Logic. In *Proc. Twenty-Fourth Annual Symposium on Foundations of Computer Science (FOCS)*, pages 195–200. IEEE, 1983.
- [Pau02] M. Pauly. A Modal Logic for Coalitional Power in Games. *Journal of Logic and Computation*, 12(1):149–166, 2002.

- [PL03] W. Penczek and A. Lomuscio. Verifying Epistemic Properties of Multi-Agent Systems via Model Checking. *Fundamenta Informaticae*, 55(2):167–185, 2003.
- [PR93] J. Pinto and R. Reiter. Temporal Reasoning in Logic Programming: A Case for the Situation Calculus. In *Proc. International Conference on Logic Programming (ICLP)*, pages 203–221, 1993.
- [PR99] F. Pirri and R. Reiter. Some Contributions to the Metatheory of the Situation Calculus. *Journal of the ACM*, 46(3):325–361, 1999.
- [PR03] M. Pauly and R. Parikh. Game Logic — An Overview. *Studia Logica*, 75(2):165–182, 2003.
- [PRS01] Artificial Intelligence Center, SRI International, Menlo Park, CA. *Procedural Reasoning System Users’s Guide*, March 13 2001. Version 2.0. <http://www.sri.com>.
- [Rao96] A. S. Rao. AgentSpeak(L): BDI Agents Speak Out in a Logical Computable Language. In *Agents Breaking Away, proc. Seventh European Workshop on Modelling Autonomous Agents in a Multi-Agent World, (MAAMAW)*, volume 1038 of *Lecture Notes in Computer Science*, pages 42–55. Springer-Verlag, 1996.
- [Rei80] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.
- [RG91] A. S. Rao and M. P. Georgeff. Modeling Rational Agents within a BDI-Architecture. In J. Allen, R. Fikes, and E. Sandewall, editors, *Proc. Second International Conference on Principles of Knowledge Representation and Reasoning: (KR)*, pages 473–484. Kaufmann, San Mateo, CA, 1991.
- [RG98] A. S. Rao and M. P. Georgeff. Decision Procedures for BDI logics. *Journal of Logic and Computation*, 8(3):293–343, 1998.
- [RG95] A. S. Rao and M. P. Georgeff. BDI Agents: From Theory to Practice. In *Proc. First International Conference on Multi-Agent Systems (ICMAS)*, San Francisco, USA, June 1995. pp312–319.
- [RvHM03] B.M. van Riemsdijk and W. van der Hoek and J.-J.Ch Meyer. Agent Programming in Dribble: from Beliefs to Goals with Plans. In *Proc. Second International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp 393–400. 2003.
- [RKJ04] A. Regayeg, A. Hadj Kacem, and M. Jmaiel. Specification and Verification of Multi-Agent Applications using Temporal Z. In *Proc. IEEE/WIC/ACM International Conference on Intelligent Agent Technology (IAT)*, pp. 260-266. Beijing, China, September, 2004.
- [RL04] F. Raimondi, A. Lomuscio. Verification of Multiagent Systems via Ordered Binary Decision Diagrams: An Algorithm and Its Implementation. In *Proc. Third International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)* pp630-637, 2004.
- [RS02] M. Ryan and P-Y. Schobbens. Agents and Roles: Refinement in Alternating-Time Temporal Logic. In Meyer and Tambe [MT02], pages 100–114.
- [Sad06] F. Sadri. Using the KGP Model of Agency to Design Applications (tutorial paper). In Toni and Torroni [TT06], pages 165–185.
- [Sat02] K. Satoh. Speculative Computation and Abduction for an Autonomous Agent. In *Proc. Ninth International Workshop on Non-Monotonic Reasoning (NMR)*, pages 191–199, 2002.
- [Sch00] K. Schild. On the Relationship Between BDI Logics and Standard Logics of Concurrency. In *Journal of Autonomous Agents and Multi-Agent Systems* 3(3):259–283, Kluwer Academic Publishers, 2000.
- [Sho93] Y. Shoham. Agent-Oriented Programming. *Artificial Intelligence*, 60(1):51–92, 1993.
- [SIIS00] K. Satoh, K. Inoue, K. Iwanuma, and C. Sakama. Speculative Computation by Abduction under Incomplete Communication Environments. In *Proc. Fourth International Conference on Multi-Agent Systems (ICMAS)*, pages 263–270. IEEE Press, 2000.
- [Sin96] M. P. Singh. Semantical Considerations on Some Primitives for Agent Specification. In *Agent Theories, Architectures, and Languages II*, volume 1037 of *Lecture Notes in Computer Science*, pp 49–64. Springer-Verlag, 1996.

- [Sin98] M. P. Singh. Agent communication languages: Rethinking the principles. *IEEE Computer*, 31(12):40–47, December 1998.
- [SLL02] S. Shapiro, Y. Lespérance, and H.J. Levesque. The Cognitive Agents Specification Language and Verification Environment for Multiagent Systems. In *Proc. First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. ACM Press, 2002.
- [Smi00] G. Smith. *The Object-Z Specification Language*. Advances in Formal Methods. Kluwer Academic Publishers, 2000.
- [SOC] Societies Of Computees (SOCS): A Computational Logic Model for the Description, Analysis and Verification of Global and Open Societies of Heterogeneous Computees. IST-2001-32530. <http://lia.deis.unibo.it/Research/SOCS/>.
- [Spi92] J. M. Spivey. *The Z Notation*. Prentice Hall, Hemel Hempstead, 2nd edition, 1992.
- [SRG99] M. P. Singh, A. S. Rao, and M. P. Georgeff. Formal Methods in DAI: Logic-Based Representation and Reasoning. In G. Weiß, editor, *Multiagent Systems—A Modern Approach to Distributed Artificial Intelligence*, chapter 8, pages 331–376. MIT Press, Cambridge, MA, 1999.
- [SRW98] M. P. Singh, A. S. Rao and M. Wooldridge, editors. *Intelligent Agents IV*, volume 1365 of *Lecture Notes in Artificial Intelligence*. Springer Verlag, 1998.
- [ST06] F. Sadri and F. Toni. A Formal Analysis of KGP Agents. In M. Fisher, W. van der Hoek, B. Konev, and A. Lisitsa, editors, *Proc. Tenth European Conference on Logics in Artificial Intelligence (JELIA)*, volume 4160 of *Lecture Notes in Computer Science*, pages 413–425. Springer-Verlag, 2006.
- [STT01] F. Sadri, F. Toni, and P. Torroni. Logic Agents, Dialogues and Negotiation: An Abductive Approach. In *Proc. AISB'01 Convention*, York, UK, March 2001.
- [STT02a] F. Sadri, F. Toni, and P. Torroni. An Abductive Logic Programming Architecture for Negotiating Agents. In *Proc. Eighth European Conference on Logics in Artificial Intelligence (JELIA)*, volume 2424 of *Lecture Notes in Computer Science*, pages 419–431. Springer-Verlag, September 2002.
- [STT02b] F. Sadri, F. Toni, and P. Torroni. Dialogues for Negotiation: Agent Varieties and Dialogue Sequences. In Meyer and Tambe [MT02], pages 405–421.
- [STT03] F. Sadri, F. Toni, and P. Torroni. Minimally Intrusive Negotiating Agents for Resource Sharing. In *Proc. Eighteenth International Joint Conference on Artificial Intelligence (IJCAI)*. Morgan Kaufmann Publishers, August 2003.
- [SY02] K. Satoh and K. Yamamoto. Speculative Computation with Multi-Agent Belief Revision. In *Proc. First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 897–904. ACM Press, 2002.
- [Tho95] S. R. Thomas. The PLACA agent programming language. In Wooldridge and Jennings [WJ95b], pages 307–319.
- [Tho00] R. Thomason. Desires and Defaults: A Framework for Planning with Inferred Goals. In *Proc. International Conference on Knowledge Representation and Reasoning (KR&R)*, pages 702–713. Morgan Kaufmann, 2000.
- [Ton06] F. Toni. Multi-agent systems in computational logic: Challenges and outcomes of the SOCS project. In Toni and Torroni [TT06], pages 420–426.
- [TT06] F. Toni and P. Torroni, editors. *Computational Logic in Multi-Agent Systems VI*, volume 3900 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, 2006.
- [VHB⁺00] W. Visser, K. Havelund, G. Brat, and S. Park. Model Checking Programs. In *Proc. Fifteenth International Conference on Automated Software Engineering (ASE)*, pages 3–12. IEEE Computer Society, 2000.
- [Wal04] C. Walton. Model checking multi-agent web services. In *Proc. 2004 Spring Symposium on Semantic Web Services*, Stanford, California, USA, March 2004.

- [Wal05] C. Walton. Model checking agent dialogues. In *Proc. Workshop on Declarative Agent Languages and Technologies (DALT)*, volume 3476 of *Lecture Notes in Artificial Intelligence*, 2005.
- [WDF98] M. Wooldridge, C. Dixon, and M. Fisher. A Tableau-Based Proof Method for Temporal Logics of Knowledge and Belief. *Journal of Applied Non-Classical Logics*, 8(3):225–258, 1998.
- [WJ95a] M. Wooldridge and N. R. Jennings. Agent Theories, Architectures, and Languages: A Survey. In Wooldridge and Jennings [WJ95b].
- [WJ95b] M. Wooldridge and N. R. Jennings, editors. *Intelligent Agents*, volume 890 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, 1995.
- [WLW⁺06] D. Walther, C. Lutz, F. Wolter and M. Wooldridge. ATL satisfiability is indeed Exptime-complete. To appear in *Journal of Logic and Computation*.
- [Woo92] M. Wooldridge. *The Logical Modelling of Computational Multi-Agent Systems*. PhD Thesis, Department of Computation, UMIST, Manchester, UK, 1992.
- [Woo98] M. Wooldridge. Verifiable semantics for agent communication languages. In *Proc. Third International Conference on Multi-Agent Systems (ICMAS)*, pages 349–365, Paris, France, 1998.
- [Woo00] M. Wooldridge. *Reasoning about Rational Agents*. The MIT Press, Cambridge, MA, 2000.
- [WPL05] B. Wozna, W. Penczek and A. Lomuscio. Bounded Model Checking for Deontic Interpreted Systems. *Electronic Notes in Theoretical Computer Science* 126:93–114, 2005.
- [WS04] F. Wan and M. P. Singh. Mapping dooley graphs and commitment causality to the π -calculus. In *Proc. Third International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 412–419, New York, NY, 2004. ACM Press.
- [W3C1] World-Wide Web Consortium. A Model-Theoretic Semantics for DAML+OIL, March 2001. <http://www.w3.org/TR/daml+oil-model>.
- [W3C2] World-Wide Web Consortium. OWL Web Ontology Language Reference [W3C Candidate Recommendation], August 2003. <http://www.w3.org/TR/owl-ref>.
- [YS02] P. Yolum and M.P. Singh. Flexible Protocol Specification and Execution: Applying Event Calculus Planning Using Commitments. In *Proc. First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 527–534. ACM Press, 2002.
- [YS04] P. Yolum and M.P. Singh. Reasoning about Commitments in the Event Calculus: An Approach for Specifying and Executing Protocols. *Annals of Mathematics and Artificial Intelligence*, 42(1-3):227–253, 2004.