

An Executable Specification of a Formal Argumentation Protocol

Alexander Artikis^a, Marek Sergot^b, Jeremy Pitt^c

^a*Institute of Informatics & Telecommunications, NCSR “Demokritos”,
Athens, 15310, Greece*

^b*Department of Computing, Imperial College London, SW7 2AZ, UK*

^c*Department of Electrical & Electronic Engineering, Imperial College London,
SW7 2BT, UK*

Abstract

We present a specification, in the action language $\mathcal{C}+$, of Brewka’s reconstruction of a theory of formal disputation originally proposed by Rescher. The focus is on the procedural aspects rather than the adequacy of this particular protocol for the conduct of debate and the resolution of disputes. The specification is structured in three separate levels, covering (i) the physical capabilities of the participant agents, (ii) the rules defining the protocol itself, specifying which actions are ‘proper’ and ‘timely’ according to the protocol and their effects on the protocol state, and (iii) the permissions, prohibitions, and obligations of the agents, and the sanctions and enforcement strategies that deal with non-compliance. Also included is a mechanism by which an agent may object to an action by another participant, and an optional ‘silence implies consent’ principle. Although comparatively simple, Brewka’s protocol is thus representative of a wide range of other more complex argumentation and dispute resolution procedures that have been proposed. Finally, we show how the ‘Causal Calculator’ implementation of $\mathcal{C}+$ can be used to animate the specification and to investigate and verify properties of the protocol.

Key words: argumentation, disputation, protocol, norm, multi-agent system, specification, action language

Email addresses: a.artikis@acm.org (Alexander Artikis), mjs@doc.ic.ac.uk (Marek Sergot), j.pitt@imperial.ac.uk (Jeremy Pitt).

1 Introduction

One of the main tasks in the formal specification and analysis of (open) multi-agent systems (MAS) is the representation of the protocols and procedures for agent interactions, and the norms of behaviour that govern these interactions. Examples include protocols for exchanging information, for negotiation, and for resolving disputes.

It has been argued that a specification of systems of this type should satisfy at least the following two requirements: first, the interactions of the members should be governed by a formal, declarative, verifiable and meaningful semantics (Singh, 2000); and second, to cater for the possibility that agent behaviour may deviate from what is prescribed, agent interactions can usefully be described in terms of permissions and obligations (Jones and Sergot, 1993).

We have been developing a theoretical framework for the executable specification of open agent systems that addresses the aforementioned requirements (Artikis et al., 2003; Artikis, 2003; Artikis et al., 2006). We adopt the perspective of an external observer, thus taking into account only externally observable behaviours and not the internal architectures of the individual agents, and view agent systems as instances of *normative systems* (Jones and Sergot, 1993) whereby constraints on agents' behaviour (or social constraints) are specified in terms of their permissions, their institutional power to effect changes and bring about certain states of affairs, and their rights and obligations to one another. We employ an action formalism to specify the social constraints governing the behaviour of the members and then use a computational framework to animate the specification and investigate its properties. For the action formalism, we have employed the Event Calculus (Kowalski and Sergot, 1986), the action language $\mathcal{C}+$ (Giunchiglia et al., 2004), and an extended form of $\mathcal{C}+$ specifically designed for modelling the institutional aspects of agent systems (Sergot, 2004a,b; Sergot and Craven, 2006).

In this paper we demonstrate how the theoretical and computational frameworks can be used with the language $\mathcal{C}+$ to specify and execute an argumentation protocol based on Brewka's reconstruction (Brewka, 2001), in the Situation Calculus (Reiter, 2001), of a theory of formal disputation originally proposed by Rescher (1977). We presented a preliminary formulation in an earlier paper (Artikis et al., 2003). This present paper is a refined and much extended version.

We are focussing here on the procedural aspects of the protocol rather than on the underlying logic of disputation employed by Brewka or on the adequacy of this particular protocol for the conduct of debate and the resolution of disputes. The features of Brewka's protocol are representative of a wide range

of other more complex argumentation and dispute resolution procedures that have been proposed in the literature, and to which the methods of this paper can be similarly applied.

The specification of the argumentation protocol is structured into three separate levels, covering:

- (i) the physical capabilities of the participant agents (in the present context, the messages/utterances each agent is actually capable of transmitting);
- (ii) the rules defining the protocol itself, specifying which actions are ‘proper’ and ‘timely’ according to the protocol and their effects on the protocol state;
- (iii) the permissions, prohibitions and obligations of the agents, and the sanctions and enforcement strategies that deal with non-compliance.

In any given implementation of the protocol, it may or may not be permitted for an agent to perform an action that is not proper or timely; conversely, there may be protocol actions that are proper and timely but that are nevertheless not permitted under certain circumstances, because, for instance, they lead to protocol runs with undesirable properties. The rules comprising level (ii) of the specification correspond to *constitutive norms* that define the meaning of the protocol actions. Levels (i) and (iii), respectively, can be seen as representing the *physical* and *normative* environment within which the protocol is executed. We have also been concerned with the concept of social role. Briefly, a role is associated with a set of (role) preconditions that agents must satisfy in order to be eligible to occupy that role and a set of (role) constraints that govern the behaviour of the agents once they occupy that role. We will not discuss role assignment in this paper. For the example in this paper, we will assume for simplicity that the participant agents are already assigned to certain roles, and that these roles do not change during an execution of the protocol.

A note on terminology In the earlier version of this paper ([Artikis et al., 2003](#)), and in the treatment of other examples, we defined a protocol by specifying the conditions under which an action was said to be ‘valid’ according to the protocol. Here, we have employed a finer structure, further classifying ‘valid’ actions as proper or timely, in line with suggestions that have also been made by Prakken et al. ([Prakken, 1998](#); [Prakken and Gordon, 1999](#)). A ‘valid’ action in our earlier terminology is one that is both proper and timely. Other terminology in common use employs the term ‘successful’ where we say ‘valid’: one then distinguishes between an action, such as an utterance or the transmission of a message of a certain form, which is an ‘attempt’ to make a claim, say, and the conditions under which the attempt to claim is ‘successful’ (sometimes, ‘effective’). We prefer to avoid the term ‘successful’ since even an unsuccessful ‘attempt’ can have effects on the protocol state. We also avoid use of the term ‘legal’ for ‘valid’ or ‘successful’ since it is ambiguous

as to whether it refers to the constitutive element of the protocol (level (ii) of our specification) or the normative environment in which the protocol is executed (level (iii)). Also related is the concept of institutional (or ‘institutionalised’) power (sometimes, ‘competence’ or ‘capacity’). This refers to the characteristic feature of institutions — legal systems, formal organisations, or informal groupings — whereby designated agents, often when acting in specific roles, are empowered to create or modify facts of special significance in that institution — *institutional facts* in the terminology of [Searle \(1969\)](#). (See e.g. ([Makinson, 1986](#); [Jones and Sergot, 1996](#)) for further discussion and references to the literature.) Thus in the present example it is natural to say that, under certain circumstances, an agent acting in a certain role has power (competence, capacity) to declare the dispute resolved in favour of one or other of the protagonists; or that in certain circumstances an agent has power to object to an action by one of the other participants; or more generally, that the argumentation protocol defines the conditions under which an agent has the power to perform one of the argumentation actions. We will not refer explicitly to power in the specification of the argumentation protocol presented here. The classification of actions into proper and timely already provides a more detailed specification.

In this paper we use the language $\mathcal{C}+$ to formulate the specification. An advantage of $\mathcal{C}+$, compared with other action formalisms, is that it can be given an explicit semantics in terms of transition systems. This enables us to analyse and prove properties of the protocol. The concluding sections of the paper present some illustrative examples.

The paper is structured as follows. First, we briefly describe the $\mathcal{C}+$ language. Second, we present the ‘Causal Calculator’ software implementation, a computational framework for executing specifications formalised in $\mathcal{C}+$. Third, we summarise Brewka’s reconstruction of Rescher’s theory of formal disputation. Fourth, we specify, prove properties of, and execute (a form of) Brewka’s argumentation protocol with the use of $\mathcal{C}+$ and the Causal Calculator. Finally, we discuss related research, summarise the presented work, and point out directions for further investigations.

2 The $\mathcal{C}+$ Language

$\mathcal{C}+$, as mentioned above, is an action language with an explicit transition systems semantics. We describe here the version of $\mathcal{C}+$ presented in ([Giunchiglia et al., 2004](#)).

2.1 Basic Definitions

A *multi-valued propositional signature* is:

- a set σ of symbols called *constants*, and
- for each constant $c \in \sigma$, a non-empty finite set $dom(c)$ of symbols, disjoint from σ , called the *domain* of c .

For simplicity, in this presentation we will assume that every domain contains at least two elements.

An *atom* of signature σ is an expression of the form $c = u$ where $c \in \sigma$ and $u \in dom(c)$. A Boolean constant is one whose domain is the set of truth values $\{\mathbf{t}, \mathbf{f}\}$. When c is a Boolean constant we often write c for $c = \mathbf{t}$ and $\neg c$ for $c = \mathbf{f}$. A *formula* φ of signature σ is any propositional combination of atoms of σ . An *interpretation* I of σ is a function that maps every constant in σ to an element of its domain. An interpretation I *satisfies* an atom $c = u$ if $I(c) = u$. The satisfaction relation is extended from atoms to formulas according to the standard truth tables for the propositional connectives. A *model* of a set X of formulas of signature σ is an interpretation of σ that satisfies all formulas in X . If every model of a set X of formulas satisfies a formula φ then X *entails* φ , written $X \models \varphi$.

2.2 Syntax

The representation of an action domain in $\mathcal{C}+$ consists of *fluent* constants and *action* constants.

- Fluent constants are symbols characterising a state. They are divided into two categories: simple fluent constants and statically determined fluent constants. Simple fluent constants are related to actions by *dynamic laws* (that is, laws describing a transition from a state s_i to its successor state s_{i+1}). Statically determined fluent constants are characterised by *static laws* (that is, laws describing an individual state) relating them to other fluent constants. Static laws can also be used to express constraints between simple fluent constants. Static and dynamic laws are defined below.
- Action constants are symbols characterising state transitions.

An *action signature* (σ^f, σ^a) is a non-empty set σ^f of fluent constants and a non-empty set σ^a of action constants. An *action description* D in $\mathcal{C}+$ is a non-empty set of *causal laws* that define a transition system of a particular type. A causal law can be either a *static law* or a *dynamic law*. A static law

is an expression

$$\text{caused } F \text{ if } G \tag{1}$$

where F and G are formulas of fluent constants. In a static law, constants in F and G are evaluated on the same state. A dynamic law is an expression

$$\text{caused } F \text{ if } G \text{ after } H \tag{2}$$

where F , G and H are formulas such that every constant occurring in F is a simple fluent constant, every constant occurring in G is a fluent constant, and H is any combination of fluent constants and action constants. In a transition from state s_i to state s_{i+1} , constants in F and in G are evaluated on s_{i+1} , fluent constants in H are evaluated on s_i and action constants in H are evaluated on the transition itself. F is called the *head* of the static law (1) and the dynamic law (2).

The full $\mathcal{C}+$ language also provides *action dynamic laws*, which are expressions of the form

$$\text{caused } \alpha \text{ if } H$$

where α is a formula containing action constants only and H is a formula of action and fluent constants. We will not use action dynamic laws in this paper and so omit the details in the interests of brevity.

The $\mathcal{C}+$ language provides various abbreviations for common forms of causal laws. For example, a dynamic law of the form

$$\text{caused } F \text{ if } \top \text{ after } H \wedge \alpha$$

where α is a formula of action constants is often abbreviated as

$$\alpha \text{ causes } F \text{ if } H$$

In the case where H is \top the above abbreviation is expressed as follows:

$$\alpha \text{ causes } F$$

When presenting the argumentation protocol specification, we will often employ the **causes** abbreviation to express the effects of the agents' actions. We will also employ the $\mathcal{C}+$ abbreviation

$$\text{default } F$$

which is shorthand for the static law

$$\text{caused } F \text{ if } F$$

expressing that F holds in the absence of information to the contrary.

When it aids readability, we will write

$$F \text{ iff } G$$

as a shorthand for the pair of static laws

$$\text{caused } F \text{ if } G$$

and

$$\text{default } \neg F$$

Two further abbreviations that we will employ are **nonexecutable** and **inertial**; a dynamic law of the form

$$\text{caused } \perp \text{ if } \top \text{ after } \alpha \wedge H$$

where α is a formula containing only action constants and H is a formula containing only fluent constants is abbreviated as:

$$\text{nonexecutable } \alpha \text{ if } H$$

In the case where H is \top the abbreviation can be written as follows:

$$\text{nonexecutable } \alpha$$

The inertia of a fluent constant c over time is represented as:

$$\text{inertial } c$$

This is an abbreviation for the *set* of dynamic laws of the form (for all values $u \in \text{dom}(c)$):

$$\text{caused } c = u \text{ if } c = u \text{ after } c = u$$

As already mentioned, a $\mathcal{C}+$ action description is a non-empty set of causal laws. Of particular interest is the sub-class of *definite* action descriptions. A $\mathcal{C}+$ action description D is *definite* if:

- the head of every causal law of D is an atom or \perp , and
- no atom is the head of infinitely many causal laws of D .

The $\mathcal{C}+$ action description in this paper will be definite.

2.3 Semantics

It is not possible in the space available here to give a full account of the $\mathcal{C}+$ language and its semantics. We trust that the $\mathcal{C}+$ language, and especially

its abbreviations, are sufficiently natural that readers can follow the presentation of the case study in later sections. Interested readers are referred to (Giunchiglia et al., 2001, 2004) for further technical details. For completeness, we summarise here the semantics of *definite* action descriptions ignoring (as we are) the presence of action dynamic laws (and assuming that the domain of every constant contains at least two elements). We emphasise the transition system semantics, as in (Sergot, 2004a).

Every action description D of $\mathcal{C}+$ defines a labelled transition system, as follows:

- States of the transition system are interpretations of the fluent constants σ^f . It is convenient to identify a state s with the set of fluent atoms satisfied by s (in other words, $s \models f = v$ if and only if $f = v \in s$ for every fluent constant f).
Let $T_{\text{static}}(s)$ denote the heads of all static laws in D whose conditions are satisfied by s :

$$T_{\text{static}}(s) =_{\text{def}} \{F \mid \text{static law (1) is in } D, s \models G\}$$

For a definite action description D , an interpretation s of σ^f is a *state* of the transition system defined by D (or simply, a state of D for short) when

$$s = T_{\text{static}}(s) \cup \text{Simple}(s)$$

where $\text{Simple}(s)$ denotes the set of simple fluent atoms satisfied by s . (So $s - \text{Simple}(s)$ is the set of statically determined fluent atoms satisfied by s .)

- Transition labels of the transition system defined by D (also referred to as *events* or *actions*) are the interpretations of the action constants σ^a .
A *transition* is a triple (s, ε, s') in which s is the initial state, s' is the resulting state, and ε is the transition label (or event). Since transition labels are interpretations of σ^a , it is meaningful to say that a transition label ε satisfies a formula α of σ^a : when $\varepsilon \models \alpha$ we sometimes say that the transition (s, ε, s') is of type α .
- Let $E(s, \varepsilon, s')$ denote the heads of all dynamic laws of D whose conditions are satisfied by the transition (s, ε, s') :

$$E(s, \varepsilon, s') =_{\text{def}} \{F \mid \text{dynamic law (2) is in } D, s' \models G, s \cup \varepsilon \models H\}$$

For a definite action description D , (s, ε, s') is a *transition of D* (or in full, a transition of the transition system defined by D) when s and s' are interpretations (set of atoms) of σ^f and ε is an interpretation of σ^a such that:

- $s = T_{\text{static}}(s) \cup \text{Simple}(s)$ (s is a state of D)
- $s' = T_{\text{static}}(s') \cup E(s, \varepsilon, s')$

For any non-negative integer m , a *path* or *history* of D of length m is a sequence

$$s_0 \varepsilon_0 s_1 \dots s_{m-1} \varepsilon_{m-1} s_m$$

where $(s_0, \varepsilon_0, s_1), \dots, (s_{m-1}, \varepsilon_{m-1}, s_m)$ are transitions of D .

3 The Causal Calculator

The Causal Calculator (CCALC) is a computational framework designed and implemented by the Action Group of the University of Texas for representing action and change in the $\mathcal{C}+$ language and performing a range of computational tasks on the resulting formalisations. CCALC has been applied to several ‘challenge problems’ (see, for example, (Akman et al., 2004; Lifschitz, 2000; Lifschitz et al., 2000)).

A detailed account of CCALC’s operation and functionality may be found in (Giunchiglia et al., 2004). This section describes the way we use *prediction*, *planning* and *postdiction* queries to execute the specification of the social constraints (or protocol rules) of an open agent system. In each type of query, CCALC has as input a definite $\mathcal{C}+$ action description D^{soc} expressing the specification of social constraints. We refer to the states of the transition system defined by D^{soc} as *social states* (or protocol states). In other words, a social state is an interpretation (with some further properties) of the fluent constants of D^{soc} . These constants express, amongst other things, which actions are physically possible, proper, timely, permitted or obligatory, and whether an agent has been sanctioned for performing forbidden actions or not complying with its obligations.

- Prediction queries. The computation of an answer to this type of query involves: (i) an initial social state expressing, amongst other things, which actions are initially physically possible, proper, timely, permitted or obligatory, and the initial sanctions of the agents (an initial social state may be partial or complete), and (ii) a *narrative*, that is, a description of temporally-sorted externally observable events of the system (a narrative is expressed as a sequence of transitions). The outcome of a prediction query (if any) is the current social state, that is, the state resulting from the events described in the narrative, expressing, amongst other things, which actions are currently physically possible, proper, timely, permitted or obligatory, and the current sanctions of the agents (see Figure 1). Computing answers to prediction queries may be performed at run-time to inform the members of a system, at any time during the execution of the system, of their permissions, obligations, sanctions and so on.
- Planning queries. Agents may issue planning queries to CCALC: (i) at design-

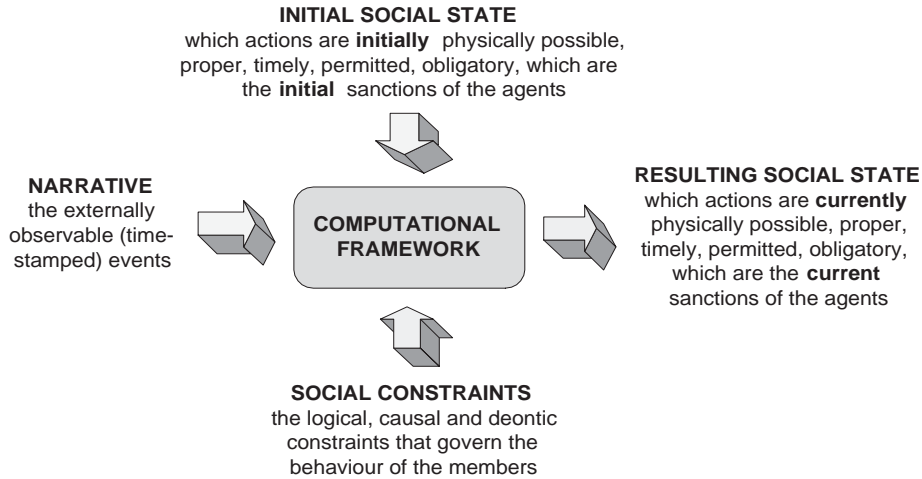


Fig. 1. Executing the Specification of an Open Agent System: Prediction Queries

time in order to generate plans that will facilitate them in avoiding run-time conflicts (say), and (ii) at run-time in order to update their plans.

- Postdiction queries. New members of a system may seek to determine the past states of that system. Similar information may be requested by agents that have ‘crashed’ and resumed their operation. Such information can be produced via the computation of answers to postdiction queries.

The computation of answers to queries may be additionally used to prove properties of the social constraints’ specification; in Section 8 we employ CICAL to prove properties of the argumentation protocol specification.

4 Rescher’s Theory of Formal Disputation

We describe, specify, prove properties of, and animate an argumentation protocol — a procedure for the resolution of a dispute — based on Brewka’s reconstruction (Brewka, 2001) of Rescher’s Theory of Formal Disputation (RTFD) (Rescher, 1977). We have picked this example because (i) in defining a set repertoire of possible moves for each participant, and their effects, it is typical of the kind of protocols that are encountered in the multi-agent systems (MAS) field, (ii) it provides a concrete and comparatively simple example of a formal procedure for the resolution of disputes, and (iii) Brewka’s formalisation in the Situation Calculus provides a natural starting point and basis for comparison. This section briefly presents RTFD. Sections 5 and 6 present Brewka’s reconstruction and our variation of RTFD respectively.

According to RTFD, argumentation may be viewed as a three-player game: the *proponent* claims a particular thesis and the *opponent* may question this thesis. The *determiner* decides whether the proponent’s thesis was successfully

defended or not. The main actions that the participants may perform are the following (p, q below are logical formulas):

- Categorical assertions. These are assertions of the form ‘ p is the case’ and are performed by the proponent.
- Cautious assertions. These are assertions of the form ‘ p is the case for all you have shown’ and are performed by the opponent.
- Provisoed assertions. These assertions are performed by either the proponent or the opponent and are expressed as follows: ‘ p generally obtains provided that q ’. A provisoed assertion of this form is accompanied with either a categorical assertion or a cautious assertion about q .

The argumentation commences with a categorical assertion by the proponent regarding the topic of the argumentation, say p . The opponent may question the topic by either a challenge of the form ‘ $\neg p$ is the case for all you have shown’ or by a provisoed denial of the form ‘ $\neg p$ generally obtains provided that q ’ and ‘ q is the case for all you have shown’. The argumentation continues in this manner until the topic has been ‘accepted’ by both the proponent and opponent (the precise meaning of ‘accepting’ a formula will be presented in the following section), in which case the determiner declares the proponent the winner, or the proponent itself does not accept the topic any more, in which case the determiner declares the opponent the winner. If neither of these alternatives take place and the participants cannot perform any additional reasonable actions, or if a deadline occurs, then the determiner decides about the winner “based on the plausibility of the proponent’s claims that were not conceded by the opponent” (Brewka, 2001, p.271). Rescher’s theory exhibits the *silence implies consent* principle. According to this principle, an agent that does not explicitly challenge a claim performed by the other is assumed to concede to the claim.

5 Brewka’s Reconstruction of RTFD

An argumentation system, according to (Brewka, 2001, Definition 4.9), includes as core components a *logic of disputation* and an *argumentation context*. In Brewka’s reconstruction, the logic of disputation is preferential default logic (Brewka and Eiter, 1998) and the argumentation context is formalised with the use of a Situation Calculus dialect (Reiter, 2001). The main actions of the protocol are the following: *claiming*, *conceding to*, *retracting*, and *denying* propositions and default rules, *declaring* the winner of the argumentation, and *objecting* to actions performed by the other participants.

The semantics of the protocol actions are given in terms of the *premises*¹ held by the proponent and opponent. The *premise*(*ag*, *q*, *s*) fluent expresses the formulas *q* that *ag* holds explicitly. The related fluent *accepts*(*ag*, *q*, *s*) is used to represent the formulas that *ag* holds implicitly: *accepts*(*ag*, *q*, *s*) expresses that *q* follows in the logic of disputation *L* from the premises explicitly held by agent *ag* in argumentation record *s*:

$$\text{accepts}(ag, q, s) \text{ iff } \{p \mid \text{premise}(ag, p, s)\} \vdash_L q \quad (3)$$

An argumentation record is a situation (in the terminology of the Situation Calculus) and so includes the history of the protocol.

The semantics of a *claim* action, for example, of a proposition or a default rule, are given by the following Situation Calculus *effect axiom*:

$$\text{premise}(ag, q, \text{do}(\text{claim}(ag, q), s)) \quad (4)$$

Expression (4) states that the successor situation following the performance of a claim action by *ag* includes a premise about the claimed proposition (or default). In expression (4) *q* represents either a proposition or a default rule of the form *n* :: *a* : *b/c* where *n* is a label associated with the default rule, *a* is the prerequisite, *b* is the justification and *c* is the consequent of the rule (Brewka, 2001).

Brewka distinguishes between ‘possible’ and ‘legal’ actions. Possible actions are specified by means of Situation Calculus *possibility axioms*. Consider the following possibility axiom of the retract action:

$$\text{poss}(\text{retract}(ag, q), s) \leftrightarrow \text{premise}(ag, q, s) \quad (5)$$

The above axiom states that it is ‘possible’ for an agent to retract a proposition (or a default) *q* if and only if that agent has a premise about that proposition (or default). The conditions that determine whether an action is possible or not are specified in a protocol-independent manner.

‘Legal’ actions, in contrast to possible actions, are specified in a protocol-dependent manner. Consider the following example of a legal action:

$$\text{legal}(\text{declare}(det, pro), s) \rightarrow \text{accepts}(pro, topic, s) \quad (6)$$

The above expression states that declaring the proponent ‘pro’ as the winner is legal only if the proponent accepts the topic of the argumentation.

¹ In work on argumentation protocols and dialogue games the term ‘commitment’ is often used where we say ‘premise’. We will not use the term ‘commitment’ here partly to keep the link with Brewka’s account, but also because ‘commitment’ has another meaning related to obligation which might cause confusion in later sections of the paper.

A point of departure of Brewka’s reconstruction from Rescher’s theory is the introduction of the ‘object’ action. The participants of an argumentation protocol may perform illegal actions; the effects of an illegal action are the same as if the action were a legal one — provided that no other participant objects to the illegal action. If some participant objects immediately (that is, no other action takes place between the illegal action and the objection), then the effects of the illegal action are ‘cancelled’.

The ‘object’ mechanism is not a new idea in the field of argumentation protocols. Prakken points out that an object mechanism of this type is part of Robert’s Rules of Order (RRO): “[t]he general rule is that anything goes until any member objects, after which RRO must be strictly applied” (Prakken, 1998, p.10). One can find similar mechanisms in most procedures for the conduct of formal debates and disputes.

Enabling agents to object to other agents’ actions can lead to a more flexible argumentation protocol. In Brewka’s modification of RTFD, for example, the proponent might choose not to object to an illegal action performed by the determiner because it (the proponent) calculates that the illegal action will serve its benefit better than having the illegal action ruled out. However, it can be argued that Brewka’s object mechanism is too simplistic, if it is a model of how argumentation processes are actually conducted, and too rigid, if it is a model of how argumentation processes ought to be conducted. Consider for example the case where an agent, say the determiner, repeatedly performs illegal actions. The proponent and opponent have to object to every illegal action performed by the determiner because if they do not object, they implicitly accept the illegal actions as legal ones. In the formalisation to be presented in the following sections, we propose a way to address this issue.

Note that according to Brewka’s treatment, an objection will ‘undo’ the effects of an illegal action if and only if the objection takes place immediately after the illegal action. If for some reason an agent fails to object immediately to an illegal action (say, another action took place between the illegal action and the objection) then it will be considered that this agent does not object and so implicitly agrees to the treatment of the illegal action as a legal one. One way of overcoming this limitation is by specifying that one may effectively object to *Act* at the latest n time-points after *Act*’s performance. Such a specification, however, would raise several complications, such as the following. Assume that *Act* takes place at time t and an agent objects to this action at t' ($t < t' < t + n$). The effects of the objection may include: (i) undoing the effects of all actions that took place between t and t' (therefore, it would be necessary to keep track of a protocol history fragment), (ii) undoing the effects of *Act*, and (iii) (possibly) applying the effects of all actions that took place between t and t' . We expect that such an object mechanism would be practical only when n is small — in any setting, computerised or not. For this reason

we will follow Brewka and assume that objections can only be made to the immediately preceding action. We will allow, however, for the possibility that more than one participant objects to the last action. Further discussion of alternative object mechanisms is beyond the scope of the paper.

Finally, Brewka’s reconstruction formalises Rescher’s ‘silence implies consent’ principle as follows: an agent, say the proponent, is assumed to have an explicit premise about a proposition, if the opponent has an explicit premise about this proposition, and as long as it (the proponent) does not deny or retract the proposition.

6 A Variation of Brewka’s Reconstruction of RTFD

Although it is our general aim in this paper to present a reconstruction of Brewka’s account of RTFD we make the following adjustments to Brewka’s version. We have in mind a setting where autonomous software agents in a MAS engage in the argumentation as part of a negotiation or dispute resolution process. In this setting, the protocol actions would be transmissions of messages. During the argumentation, the proponent, opponent and determiner must perform their chosen actions by specified *deadlines* (timeouts). Without this feature there is no practical way of controlling the exchanges, of determining whether a participant has ‘spoken’, because otherwise one might have to wait indefinitely for messages to arrive over the communication channels. For similar reasons, it is also necessary to impose a limit on the number of exchanges, or on the total elapsed time for the argumentation process. (Although Brewka states that the argumentation regarding a proposition may terminate due to a deadline, he provides no further details about how this would work and how such deadlines would affect the argumentation process.) Having introduced deadlines, it is necessary to express the conditions in which an action can be said to be timely, that is, whether a participant has ‘spoken’ by the specified deadline, and so on.

We further refine Brewka’s distinction between ‘possible’ and ‘legal’ actions. We consider not only what kinds of actions are proper (one possible interpretation of the Situation Calculus predicate *poss* — see, for example, axiom (5)) but also which of these actions each agent will be *practically able* to perform at each stage of a given implementation. Moreover, we express the conditions in which an action can be said to be permitted (one possible interpretation of Brewka’s ‘legal’ actions) or even obligatory. We specify enforcement strategies to deal with the performance of forbidden actions and non-compliance with obligations.

Finally, even though we depart from Brewka’s reconstruction of RTFD in the

aforementioned points, we maintain (a form of) the object mechanism.

Here then is our variant of Brewka’s reconstruction of RTFD. We will refer to it as RTFD*. The argumentation commences when the proponent claims the topic of the argumentation — any other action does not count as the commencement of the protocol. The protagonists (proponent and opponent) then take it in turn to perform actions. Each turn lasts for a specified time period during which the protagonist may perform several actions (send several messages) up to some specified limit. After each such action the other participants are given an opportunity to object within another specified time period. In other words, *Ag*’s action *Act* is followed by a time period during which *Ag* may not perform any actions and the other participants may object to *Act*. The determiner may declare the winner only at the end of the argumentation, that is, when the specified period for the argumentation elapses. (Other specifications of the protocol’s terminating conditions are possible — for instance, the argumentation may end earlier than the specified time period if a protagonist concedes to the other’s arguments.) If at the end of the argumentation both the proponent and opponent have accepted the topic of the argumentation, then the determiner may only declare the proponent the winner. If, however, the proponent does not accept the topic then the determiner may only declare the opponent the winner. Finally, if the proponent accepts the topic and the opponent does not, the determiner has *discretion* to declare either of them the winner. It may also have an *obligation* to decide one way or the other, depending on which version of the protocol we choose to adopt.

7 Specifying RTFD*

We present a $\mathcal{C}+$ action description D^{RTFD^*} that expresses the specification of RTFD*. Table 1 shows a subset of the action signature (σ^f, σ^a) of D^{RTFD^*} . Variables are written with an upper-case first letter and constants start with a lower-case letter. The intended reading of the constants of (σ^f, σ^a) will be explained during the presentation of the RTFD* specification.

There are three roles in the argumentation protocol: proponent, opponent, and determiner. Although our specification does not rely on the assumption that there is at most one agent occupying a role in any given execution of the protocol, in the concrete example presented we will deal with the usual case where there are three agents, one in each role. In the protocol presented agents do not change role. We will call *pro* the agent that occupies the role of the proponent, *opp* the agent that occupies the role of the opponent and *det* the agent that occupies the role of the determiner. Accordingly, the example

Table 1
A Subset of the Action Signature of D^{RTFD^*}

Variables:	Domain:
Ag, Ag'	$\{pro, opp, det\}$
$Protag, Protag'$	$\{pro, opp\}$
Det	$\{det\}$
P, Q	a finite set of propositions/default rules
Act	$\{claim(Protag, Q), concede(Protag, Q), retract(Protag, Q),$ $deny(Protag, Q), declare(Det, Protag)\}$
Rigid Constants:	Domain:
$role_of(Ag)$	$\{proponent, opponent, determiner\}$
$topic$	a finite set of propositions
$implies(P, Q)$	Boolean
Simple Fluent Constants:	Domain:
$turn$	$\{proponent, opponent, determiner\}$
$premise(Protag, Q)$	$\{t, f, u\}$
$initialState, sanctioned(Ag)$	Boolean
$winner$	$\{pro, opp, none\}$
Statically Determined Fluent Constants:	Domain:
$proper(Act), timely(Act), per(Act), obl(Act),$ $objectionable(Act), accepts(Protag, Q), fair$	Boolean
$winning$	$\{pro, opp, none\}$
Action Constants σ^{act} (Boolean):	
$claim(Protag, Q), concede(Protag, Q), retract(Protag, Q),$ $deny(Protag, Q), declare(Det, Protag)$	
Action Constants σ^{obj} (Boolean):	
$objected(Ag)$	

action description contains the following three static laws²:

$$\begin{aligned}
 role_of(pro) &= proponent \text{ if } \top \\
 role_of(opp) &= opponent \text{ if } \top \\
 role_of(det) &= determiner \text{ if } \top
 \end{aligned}$$

The $role_of(Ag)$ fluent constants are thus ‘rigid’, in the sense that their values are the same in all states (the Ag variable ranges over the participants pro , opp and det). In other versions of the protocol, we might introduce other roles, for instance that of an ‘observer’ who does not participate in the argumentation

² For brevity, we will omit in the remaining of the paper the keyword **caused** which appears at the beginning of static and dynamic laws in the original presentation of $\mathcal{C}+$ (Giunchiglia et al., 2004).

proper but could be allowed to object to actions made by the protagonists. We do not show that variation here.

Other rigid constants are *topic*, whose value is a proposition expressing the topic of the argumentation, and Boolean fluent constants $implies(P, Q)$, used to represent the underlying logic of disputation, as described in a following section.

Each simple fluent constant of D^{RTFD^*} is inertial, that is to say, its value persists by default from one state to the next. The constraint that a fluent constant f is inertial is expressed in $\mathcal{C}+$ by means of the causal law abbreviation

$$\text{inertial } f \quad (7)$$

$\mathcal{C}+$ abbreviations were presented in Section 2.

The action constants σ^a of D^{RTFD^*} are partitioned into two sets, σ^{act} and σ^{obj} . The set σ^{act} includes the main actions (*claim*, *concede*, *retract*, *deny*, *declare*) of the argumentation, as summarised in Table 1. For convenience, we specify that exactly one action from σ^{act} takes place at each state transition. This is done by means of the following $\mathcal{C}+$ laws:

$$\text{nonexecutable } \alpha_i \wedge \alpha_j, \quad \text{for all } \alpha_i, \alpha_j \in \sigma^{\text{act}}, \alpha_i \neq \alpha_j \quad (8)$$

together with

$$\text{nonexecutable } \neg\alpha_1 \wedge \neg\alpha_2 \wedge \dots \wedge \neg\alpha_n \quad (9)$$

where $\alpha_1, \alpha_2, \dots, \alpha_n$ are the action constants of σ^{act} . The restriction expressed by laws (8) and (9) simplifies the $RTFD^*$ specification, and is also convenient when analysing executions of the protocol (see Sections 8 and 9).

The action constants of σ^{obj} are Boolean constants of the form $objected(Ag)$ used to represent that an objection has been made by agent Ag . In the specification presented in this paper we have abstracted away details of how an objection is transmitted within the specified deadline (recall that every action Act is followed by a time period during which no action may take place apart from an objection to Act). Instead, every transition of the transition system defined by D^{RTFD^*} corresponds to a *claim*, *concede*, *retract*, *deny*, or *declare* action by one of the participants together with an indication of whether that action was objected to by one or more of the others. For example, a transition satisfying the action formula

$$\text{claim}(pro, Q) \wedge \text{objected}(opp) \wedge \text{objected}(det)$$

represents a transition in which a claim that Q by pro is followed by objections

by both *opp* and *det*. Similarly, a transition satisfying

$$claim(pro, Q) \wedge \neg objected(opp) \wedge \neg objected(det)$$

expresses a claim that has not been objected to. We could have produced a more elaborate $\mathcal{C}+$ action description expressing the objection mechanism and its associated deadlines but the details are rather fiddly and we do not present them here so as not to distract attention from the main points of the paper. Notice that because of laws (8) and (9), the *objected*(*Ag*) action constants do not need to specify which action has been objected to since there is exactly one such action at each state transition. We do find it useful to specify the agent who objected, and to allow for the possibility that more than one agent objected within the deadline. Other versions of the specification may be constructed that do not include laws (8) and (9) but the representation of objections is then more intricate. We omit the details.

We begin by specifying the well-formed actions of the RTFD* protocol. A *claim*(*Ag*, *Q*) action, for example, is well-formed only if *Ag* is an agent occupying the role of the proponent or that of the opponent. There are two ways of specifying the well-formed actions. The simpler method, and the one we choose here, is to specify this in the action signature. For instance, the action signature summarised in Table 1 contains action constants of the form *claim*(*Protag*, *Q*) where *Protag* ranges over the two protagonists, *pro* and *opp*. The alternative method is to include ill-formed actions in the action signature and specify explicitly that they are ‘non-executable’, as in the following example:

$$\text{nonexecutable } claim(Ag, Q) \text{ if } role_of(Ag) = determiner$$

According to the above law, there is no transition in the transition system defined by D^{RTFD^*} corresponding to a claim made by an agent occupying the role of determiner. This latter way of dealing with ill-formed actions results in a more generic and flexible specification that would require fewer changes to accommodate participants’ changing roles (assuming protocol versions in which this was possible), new agents participating in the protocol, and so on. The former way (which we choose here) results in a specification with simpler laws (they have fewer conditions) and is thus easier to follow.

An objection is not well-formed if an agent objects to its own actions. Accordingly, we include laws of the following form

$$\text{nonexecutable } claim(Protag, Q) \wedge objected(Protag) \tag{10}$$

and similarly for the other action constants in σ^{act} . We do not support objections to objections since they are of little practical or theoretical interest.

We now present a number of causal laws expressing when a well-formed pro-

TOCOL action of RTFD* is physically (practically) possible, proper, timely, permitted or obligatory, and what the effects of an action are.

7.1 Physical Capability

The *system events* of the RTFD* specification are the timeouts — these are issued by a global clock. (To avoid clutter, several constants of the action signature of D^{RTFD^*} , including those representing timeouts, are not presented in Table 1.) A type of timeout event is used to denote the turn of each participant. When RTFD* commences (this happens when the proponent claims the topic of the argumentation) a global clock starts ‘ticking’. The first timeout signals the end of the proponent’s turn and the beginning of the opponent’s turn to ‘speak’, by setting $turn = opponent$. The next timeout signals the end of the opponent’s turn and the beginning of the proponent’s turn, by setting $turn = proponent$, and so on.

The remaining actions of the RTFD* specification are those performed by the protocol participants. We have chosen to specify that any protagonist is always capable of signalling a claim, concede, retract, deny, and object action, and the determiner is always capable of signalling a declare and object action (except that no agent is capable of objecting to its own actions). The effects of these actions are presented next.

At the initial state of the protocol the protagonists have no premises, that is, the value of every $premise(Protag, Q)$ fluent constant is **f**. The variable Q ranges over the formulas (propositions and default rules) that the two protagonists may claim, concede to, retract, and deny (see Table 1). We assume that a finite number of such formulas can be identified and specified at the outset. This is necessary for implementation in $\mathcal{C}+$ and the CCALC system, though not necessarily in other formalisms.

The protocol commences with the proponent’s claim of the topic. The effects of a claim action are expressed as follows:

$$\begin{aligned} &claim(Protag, Q) \text{ causes } premise(Protag, Q) = \mathbf{t} \text{ if} \\ &premise(Protag, Q) = \mathbf{f} \wedge \\ &\neg objected(pro) \wedge \neg objected(opp) \wedge \neg objected(det) \end{aligned} \tag{11}$$

The above expression is an abbreviation for the $\mathcal{C}+$ fluent dynamic law:

$$\begin{aligned} &premise(Protag, Q) = \mathbf{t} \text{ if } \top \text{ after} \\ &claim(Protag, Q) \wedge premise(Protag, Q) = \mathbf{f} \wedge \\ &\neg objected(pro) \wedge \neg objected(opp) \wedge \neg objected(det) \end{aligned}$$

Law (11) expresses that *Protag*’s claim of Q leads from a state in which *Protag*

has no explicit premise that Q (that is, $\text{premise}(\text{Protag}, Q) = \text{f}$) to a state in which it does have an explicit premise that Q (that is, $\text{premise}(\text{Protag}, Q) = \text{t}$), on condition that no (other) agent, pro , opp , or det , objects to the claim. An objection is only effective in blocking the effects of the claim action if it is well-founded (in a sense to be specified below). If the objection is not well-founded then it does not block the effects of the claim action (though it might have other effects, such as exposing the objecting agent to sanctions). We therefore add to law (11) the further constraint that:

$$\begin{aligned} \text{claim}(\text{Protag}, Q) \text{ causes } \text{premise}(\text{Protag}, Q) = \text{t} \text{ if} \\ \text{premise}(\text{Protag}, Q) = \text{f} \wedge \\ \neg \text{objectionable}(\text{claim}(\text{Protag}, Q)) \end{aligned} \quad (12)$$

Boolean fluent constants $\text{objectionable}(\text{Act})$ are used to represent that an objection to Act is well-founded.

When is an objection to an action Act ‘well-founded’, that is to say, when is $\text{objectionable}(\text{Act})$ true? In general, whenever Act is not a proper and timely action of the protocol. However, we find it adds flexibility to define objectionable separately. Not all improper or untimely actions need to be objectionable. As we explain in the next section, sometimes improper or untimely actions have no effect on the state of the argumentation, and in those circumstances the protocol can be simplified by ignoring objections to them. Specification of the proper, timely, and objectionable actions in RTFD* will be given in the next section.

For convenience we introduce a special abbreviation, writing

$$\text{Act p_causes } F \text{ if } G$$

(for ‘provisionally causes’) as shorthand for the pair of causal laws of the form

$$\begin{aligned} \text{Act causes } F \text{ if } G \wedge \neg \text{objected}(Ag_1) \wedge \dots \wedge \neg \text{objected}(Ag_n) \\ \text{Act causes } F \text{ if } G \wedge \neg \text{objectionable}(\text{Act}) \end{aligned}$$

where Ag_1, \dots, Ag_n are the participants in the argumentation (pro , opp , and det in the present example). When G is \top we write

$$\text{Act p_causes } F$$

The laws (11) and (12) stating the effects of a claim can thus be written succinctly as follows:

$$\begin{aligned} \text{claim}(\text{Protag}, Q) \text{ p_causes } \text{premise}(\text{Protag}, Q) = \text{t} \text{ if} \\ \text{premise}(\text{Protag}, Q) = \text{f} \end{aligned} \quad (13)$$

Suppose that protagonist Protag claims a proposition Q . Opponent Protag'

may respond to *Protag*'s claim by conceding to, or denying the claim. If *Protag'* does neither then we say that *Protag'* has an 'unconfirmed' premise that Q , denoted by $premise(Protag', Q) = u$. The value of a *premise* fluent constant is set to 'unconfirmed' as follows, for every pair of distinct protagonists *Protag* and *Protag'*:

$$\begin{aligned} &claim(Protag, Q) \text{ p_causes } premise(Protag', Q) = u \text{ if} \\ &\quad premise(Protag, Q) = f \wedge \\ &\quad premise(Protag', Q) = f \end{aligned} \tag{14}$$

In other words, *Protag*'s claim of Q leads (subject to possible objections) to a state in which *Protag'* has an unconfirmed premise that Q , provided that *Protag'* does not already have a premise that Q (that is, provided that $premise(Protag', Q) = f$). If *Protag'* already has a premise that Q (that is, $premise(Protag', Q) = t$) then its premise does not become unconfirmed, and it does not need to respond to *Protag*'s claim.

Responding to a claim, that is, conceding to or denying a claim, can be expressed as follows:

$$\begin{aligned} &concede(Protag, Q) \text{ p_causes } premise(Protag, Q) = t \text{ if} \\ &\quad premise(Protag, Q) = u \end{aligned} \tag{15}$$

$$\begin{aligned} &deny(Protag, Q) \text{ p_causes } premise(Protag, Q) = f \text{ if} \\ &\quad premise(Protag, Q) = u \end{aligned} \tag{16}$$

A claim may be retracted. The effects of a retraction are twofold:

$$\begin{aligned} &retract(Protag, Q) \text{ p_causes } premise(Protag, Q) = f \text{ if} \\ &\quad premise(Protag, Q) = t \end{aligned} \tag{17}$$

$$\begin{aligned} &retract(Protag, Q) \text{ p_causes } premise(Protag', Q) = f \text{ if} \\ &\quad premise(Protag, Q) = t \wedge \\ &\quad premise(Protag', Q) = u \end{aligned} \tag{18}$$

for every pair of distinct protagonists *Protag* and *Protag'*. According to law (17), *Protag*'s retraction of its claim that Q results in $premise(Protag, Q) = f$; according to law (18) this retraction also removes the unconfirmed premise held by the other protagonist, *Protag'*. If *Protag'* responded to *Protag*'s claim before *Protag* retracted it then the value of $premise(Protag', Q)$ would not be unconfirmed, and in that case *Protag*'s retraction will not affect the premise held by *Protag'*.

The effects of laws (13)–(18) for a single proposition Q , and assuming no objections, are summarised in Figure 2. Notice that, because the *premise* fluent constants are inertial, *claim*, *concede*, *deny*, and *retract* actions have no effect on the argumentation state if the pre-conditions in laws (13)–(18) are

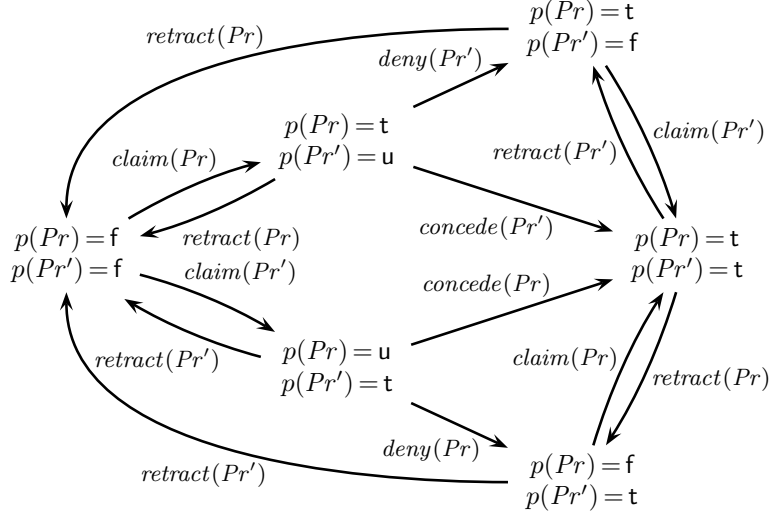


Fig. 2. Effects of *claim*, *concede*, *deny*, *retract* actions. The proposition Q is omitted for clarity. Pr , Pr' are shorthands for *Protag*, *Protag'* respectively. Similarly, $p(Pr)$, $p(Pr')$ are shorthands for *premise(Protag, Q)* and *premise(Protag', Q)*. Actions that do not change the protocol state are not shown in the diagram.

not satisfied. Transitions corresponding to such actions are omitted from the diagram for clarity.

At the end of the argumentation the determiner may declare the winner — the effects of this action can be expressed as follows:

$$\text{declare}(\text{Det}, \text{Protag}) \text{ p_causes } \text{winner} = \text{Protag} \quad (19)$$

Det represents the agent occupying the role of the determiner. The fluent constant $\text{winner} = \text{Protag}$ expresses that the winner of the argumentation is *Protag*. The use of *p_causes* here again deals with the possibility that some agent other than Det objects to the declaration.

Laws (13)–(19) express the effects of the main protocol actions. Suppose that a protagonist claims a proposition Q when it already has a premise that Q , or that it sends a *retract(Protag, Q)* message when it has no premise that Q . Such actions are not proper, according to the protocol. Moreover, an action may be untimely, as when, for instance, a protagonist speaks out of turn. Some improper or untimely actions have no effects on the state of the argumentation. Others do have effects — however, their effects may be blocked by objections. In the next section we specify when a protocol action is proper and timely, and we specify when an action is objectionable.

Table 2
 Proper and Timely Actions in D^{RTFD^*}

Action	proper	timely
$claim(Protag, Q)$	$premise(Protag, Q) = \mathbf{f}$	$(\neg initialState \vee topic = Q) \wedge turn = role_of(Protag)$
$concede(Protag, Q)$	$premise(Protag, Q) = \mathbf{u}$	$\neg initialState \wedge turn = role_of(Protag)$
$retract(Protag, Q)$	$premise(Protag, Q) = \mathbf{t}$	$\neg initialState \wedge turn = role_of(Protag)$
$deny(Protag, Q)$	$premise(Protag, Q) = \mathbf{u}$	$\neg initialState \wedge turn = role_of(Protag)$
$declare(Det, Protag)$	$winner \neq Protag \wedge (winning = Protag \vee winning = none)$	$turn = role_of(Det)$

7.2 Proper and Timely Actions

The second column of Table 2 shows the conditions in which the main protocol actions are said to be proper. A $claim(Protag, Q)$ action, for example, is proper if and only if $Protag$ does not have a premise that Q , that is, $premise(Protag, Q) = \mathbf{f}$:

$$proper(claim(Protag, Q)) \text{ iff } premise(Protag, Q) = \mathbf{f} \quad (20)$$

In other words, a $claim(Protag, Q)$ action is improper when:

- $premise(Protag, Q) = \mathbf{t}$; it is improper to make repeated claims, that is, claim something one has a premise about.
- $premise(Protag, Q) = \mathbf{u}$; this signifies that the other protagonist $Protag'$ has claimed Q (see law (14)). A proper response to the other's claim is either a concession or a denial (see Table 2).

A $retract(Protag, Q)$ action is proper if and only if $premise(Protag, Q) = \mathbf{t}$; otherwise there is nothing to retract.

The reader may have noticed that the conditions for proper $claim$, $concede$, $deny$, and $retract$ actions coincide with the $premise(Protag, Q)$ pre-conditions in the laws (13)–(18). This is a feature of Brewka's version of RTFD (or rather, our chosen formulation of some of its details) and will not necessarily be the case in other examples. In other argumentation and dialogue protocols it is common, for example, to say that a claim by $Protag$ is not proper if it is inconsistent, in the underlying logic of disputation, with the premises currently held by $Protag$. This is easily added in our formulation but is not a feature

of Brewka’s RTFD. Other formulations of proper actions will be discussed in Section 10.

We now turn to declarations. A declaration is proper when the following conditions are satisfied:

$$\begin{aligned} \text{proper}(\text{declare}(\text{Det}, \text{Protag})) \text{ iff} \\ \text{winner} \neq \text{Protag} \wedge \\ (\text{winning} = \text{Protag} \vee \text{winning} = \text{none}) \end{aligned} \quad (21)$$

In words: declaring *Protag* the winner of the argumentation is proper if and only if:

- *Protag* has not already been declared winner, and
- *Protag* is currently ‘winning’, that is, the conflict is resolved in favour of *Protag*, or no protagonist is ‘winning’, that is, the conflict is unresolved.

The statically defined fluent constant *winning* is defined as follows:

$$\begin{aligned} \text{winning} = \text{Protag} \text{ if} \\ \text{role_of}(\text{Protag}) = \text{proponent} \wedge \\ \text{role_of}(\text{Protag}') = \text{opponent} \wedge \\ \text{topic} = Q \wedge \\ \text{accepts}(\text{Protag}, Q) \wedge \\ \text{accepts}(\text{Protag}', Q) \end{aligned} \quad (22)$$

$$\begin{aligned} \text{winning} = \text{Protag}' \text{ if} \\ \text{role_of}(\text{Protag}) = \text{proponent} \wedge \\ \text{role_of}(\text{Protag}') = \text{opponent} \wedge \\ \text{topic} = Q \wedge \\ \neg \text{accepts}(\text{Protag}, Q) \end{aligned} \quad (23)$$

$$\text{default } \text{winning} = \text{none} \quad (24)$$

Laws (22) and (23) express a resolved dispute. An *accepts(Protag, Q)* fluent constant states that *Q* follows, in the logic of disputation, from *Protag*’s explicit premises. (A discussion of the logic of disputation is presented in Section 7.5.) Law (22) deals with the case that both protagonists accept the topic *Q* of the dispute; law (23) deals with the case where the proponent no longer accepts the topic *Q*. Law (24) corresponds to an unresolved dispute; in the absence of information to the contrary (that is, laws (22) and (23)), the dispute is unresolved.

Even if proper, an action may not be in compliance with the protocol specification — for example, it may be untimely. The third column of Table 2 shows the conditions in which the main protocol actions are said to be timely. At the initial state of the protocol, expressed by the simple fluent constant *initialState*, only the proponent’s claim of the argumentation topic is timely

Table 3
 Objectionable Actions in D^{RTFD^*}

Action	objectionable
$claim(Protag, Q)$	$proper(claim(Protag, Q)) \wedge \neg timely(claim(Protag, Q))$
$concede(Protag, Q)$	$proper(concede(Protag, Q)) \wedge \neg timely(concede(Protag, Q))$
$retract(Protag, Q)$	$proper(retract(Protag, Q)) \wedge \neg timely(retract(Protag, Q))$
$deny(Protag, Q)$	$proper(deny(Protag, Q)) \wedge \neg timely(deny(Protag, Q))$
$declare(Det, Protag)$	$\neg(proper(declare(Det, Protag)) \wedge timely(declare(Det, Protag))) \wedge winner \neq Protag$

(at the initial protocol state $turn = proponent$). At the other protocol states a protagonist's claim, concession, retraction, or denial is timely if it is the protagonist's turn to 'speak'. A timely claim, for instance, is expressed in $\mathcal{C}+$ as follows:

$$\begin{aligned}
 timely(claim(Protag, Q)) \text{ iff} \\
 (\neg initialState \vee topic = Q) \wedge \\
 turn = role_of(Protag)
 \end{aligned} \tag{25}$$

Declarations are defined to be timely as follows:

$$timely(declare(Det, Protag)) \text{ iff } turn = role_of(Det) \tag{26}$$

A declaration is timely if and only if it is the determiner's turn to 'speak' — this happens when the argumentation time elapses.

In this version we do not allow for early declarations, in the case where both protagonists accept the argumentation topic, or the proponent does not accept it, before the argumentation time elapses; this can be easily adjusted.

What about improper declarations? Or untimely actions? Some of these actions do have effects on the protocol state — however, these effects may be blocked by a well-founded objection. Table 3 presents the conditions in which an objection to Act is well-founded, that is, when Act is objectionable. In the present example a claim is objectionable if and only if it is untimely and proper:

$$\begin{aligned}
 objectionable(claim(Protag, Q)) \text{ iff} \\
 proper(claim(Protag, Q)) \wedge \neg timely(claim(Protag, Q))
 \end{aligned} \tag{27}$$

According to the protocol specified above, improper claims do not have effects on the protocol state. It is therefore not necessary to object to them. Similarly, a concession, retraction or denial is objectionable if and only if it is untimely and proper.

A $declare(Det, Protag)$ action is objectionable if and only if: (i) it is not proper and timely, and (ii) $winner \neq Protag$ (see Table 3). Notice that if

$winner = Protag$ at the time of a $declare(Det, Protag)$ action then this action will not have any effects on the value of the $winner$ fluent constant. Accordingly, we choose to say that an improper or untimely $declare(Det, Protag)$ action is objectionable only when $winner \neq Protag$ at the time of the declaration.

We presented the conditions in which an action is considered proper or timely, and the consequences of the performance a proper/improper, timely/untimely action. In the following section we focus on a separate issue: the circumstances in which a protocol participant is permitted or even obliged to perform a particular action. In Section 7.4 we discuss potential consequences of the performance of forbidden actions and non-compliance with obligations.

7.3 Permitted Actions

In Section 5 we identified a side-effect of the object mechanism: Ag 's repeated performance of objectionable actions requires repeated objections from some other participant Ag' otherwise the effects of Ag 's objectionable actions are not blocked. One way of addressing this issue is to say that the performance of repeated objectionable actions is not *permitted*. Consider the following example specification of permission:

$$\begin{aligned} per(concede(Protag, Q)) \text{ iff} \\ \neg objectionable(concede(Protag, Q)) \vee \\ objActions(Protag) < 10 \end{aligned} \quad (28)$$

Here, there is an upper limit (10 in the example) on the total number of objectionable actions a protagonist $Protag$ is permitted to make. The *per* fluent constants express permission; the simple fluent constants $objActions(Ag)$ record the number of objectionable actions issued by Ag . The formulation of $\mathcal{C}+$ laws to maintain the $objActions(Ag)$ fluent constants is very straightforward and we omit the details. Laws similar to (28) express the permission to claim, retract, deny and declare (we may impose a lower threshold for the permitted objectionable declarations). In the present example, an objection is never objectionable; therefore, an objection is always permitted.

In another formulation, $objActions(Ag)$ could express the number of objectionable actions issued by Ag to which well-founded (successful) objections have been made by other participants. In this case, Ag would be permitted to perform objectionable actions until the number of successful objections issued by the other participants to Ag 's actions has reached the specified limit.

We could formulate further permission laws to classify as not permitted other kinds of exchanges, such as repeated claims and retractions by the protagonists

(‘yes it is’; ‘no it isn’t’; ‘yes it is’; ‘no it isn’t’; . . .). Although such an exchange could be both proper and timely according to the protocol, it may nevertheless be undesirable – it does not progress the resolution of the dispute.

In some cases it is meaningful to say that the determiner is not only permitted, but obliged to declare a protagonist the winner of the dispute. Such an obligation will arise when the following conditions hold:

$$\begin{aligned}
obl(declare(Det, Protag)) \text{ iff} \\
turn = role_of(Det) \wedge \\
winner \neq Protag \wedge \\
winning = Protag
\end{aligned} \tag{29}$$

(The *obl* fluent constants express obligation.) According to law (29), it is obligatory to declare *Protag* the winner of the dispute if and only if: (i) it is the determiner’s turn to ‘speak’, (ii) *Protag* has not already been declared the winner, and (iii) *Protag* is currently ‘winning’ the dispute, that is, the dispute is resolved in favour of *Protag*.

Given expression (29), it is necessary to represent explicitly the relationship between the permission and obligation to declare the winner. For instance, the protocol should never reach a state in which both $obl(declare(Det, Protag))$ and $per(declare(Det, Protag'))$ are true for $Protag \neq Protag'$. One way is to add the following laws, for every pair of distinct protagonists *Protag* and *Protag'*:

$$per(declare(Det, Protag)) \text{ if } obl(declare(Det, Protag)) \tag{30}$$

$$per(declare(Det, Protag)) \text{ if } \neg obl(declare(Det, Protag')) \wedge \neg objectionable(declare(Det, Protag)) \tag{31}$$

$$per(declare(Det, Protag)) \text{ if } \neg obl(declare(Det, Protag')) \wedge objActions(Det) < 10 \tag{32}$$

$$\text{default } \neg per(declare(Det, Protag)) \tag{33}$$

According to laws (30)–(33), if the determiner is obliged to declare *Protag* the winner then it is also permitted to declare *Protag*, but not *Protag'*. When there is no obligation on the determiner, at most ten objectionable declarations are permitted (that is, the permission to declare is expressed as the permission to perform any other protocol action).

It is not meaningful to associate obligations with the remaining protocol actions; therefore, we do not need to update the specification of permitted claims, concessions, retractions, denials or objections. Table 4 presents the conditions in which a main protocol action is permitted or obligatory. Clearly, different

Table 4
Permission and Obligation in D^{RTFD^*}

Action	per	obl
$claim(Protag, Q)$	$\neg objectionable(claim(Protag, Q)) \vee$ $objActions(Protag) < 10$	\perp
$concede(Protag, Q)$	$\neg objectionable(concede(Protag, Q)) \vee$ $objActions(Protag) < 10$	\perp
$retract(Protag, Q)$	$\neg objectionable(retract(Protag, Q)) \vee$ $objActions(Protag) < 10$	\perp
$deny(Protag, Q)$	$\neg objectionable(deny(Protag, Q)) \vee$ $objActions(Protag) < 10$	\perp
$declare(Det, Protag)$	$obl(declare(Det, Protag) \vee$ $(\neg obl(declare(Det, Protag'))) \wedge$ $(\neg objectionable(declare(Det, Protag)))$ \vee $objActions(Det) < 10)$	$turn = role_of(Det) \wedge$ $winner \neq Protag \wedge$ $winning = Protag$

specifications are possible. One may forbid, for instance, repetitive improper claims, concessions, retractions, and denials, and ill-founded objections (objections to *Act* when *Act* is not objectionable). Although these actions have no effects on the protocol state and thus are not objectionable, they may nevertheless be forbidden (not permitted) in order to deter participants from performing ‘meaningless’ actions, to decrease network traffic. In (Artikis et al., 2003) we presented an RTFD* specification in which all retractions were forbidden, whether proper/improper, timely/untimely, objectionable/non-objectionable. Furthermore, the obligation to declare the winner could arise even in the case of an unresolved dispute.

Note that it is practically possible for an agent to perform forbidden actions and not comply with its obligations. We introduce enforcement strategies as a way of dealing with this type of behaviour. Such strategies are presented next.

7.4 Enforcement Strategies

We want to reduce or eliminate:

- The protagonists’ performance of forbidden actions, repeated objectionable actions in this example.
- The determiner’s non-compliance with the obligation to declare the winner of the dispute. (To simplify the presentation, we will not address the determiner’s forbidden declarations.)

We employ the simple fluent constants *sanctioned* to identify the aforementioned types of behaviour. *Protag*'s forbidden claim, for example, results in initiating *sanctioned(Protag)*:

$$\begin{aligned} \text{claim}(\text{Protag}, Q) \text{ causes } \text{sanctioned}(\text{Protag}) \text{ if} \\ \neg \text{per}(\text{claim}(\text{Protag}, Q)) \end{aligned} \quad (34)$$

One way of reducing the performance of forbidden actions and non-compliance with obligations is by penalising such behaviour. For example: if at the close of the argumentation the dispute is unresolved, we could say that declaring *Protag* the winner is objectionable if *Protag* performed some forbidden actions but *Protag'* did not. In this case, the penalty *Protag* pays for its forbidden actions is that it will not win an unresolved dispute if *Protag'*: (i) does not perform forbidden actions, and (ii) objects to a potential *declare(Det, Protag)* action. (This does not necessarily imply that *Protag'* will win the dispute.) The specification of this example sanction may be expressed by updating the definition of objectionable declarations as follows (for every pair of distinct protagonists *Protag* and *Protag'*):

$$\begin{aligned} \text{objectionable}(\text{declare}(\text{Det}, \text{Protag})) \text{ if} \\ \neg \text{proper}(\text{declare}(\text{Det}, \text{Protag})) \wedge \\ \text{winner} \neq \text{Protag} \end{aligned} \quad (35)$$

$$\begin{aligned} \text{objectionable}(\text{declare}(\text{Det}, \text{Protag})) \text{ if} \\ \neg \text{timely}(\text{declare}(\text{Det}, \text{Protag})) \wedge \\ \text{winner} \neq \text{Protag} \end{aligned} \quad (36)$$

$$\begin{aligned} \text{objectionable}(\text{declare}(\text{Det}, \text{Protag})) \text{ if} \\ \text{winning} = \text{none} \wedge \\ \text{winner} \neq \text{Protag} \wedge \\ \text{sanctioned}(\text{Protag}) \wedge \\ \neg \text{sanctioned}(\text{Protag}') \end{aligned} \quad (37)$$

$$\text{default } \neg \text{objectionable}(\text{declare}(\text{Det}, \text{Protag})) \quad (38)$$

Laws (35), (36) and (38) express the definition of objectionable declarations presented in Table 3; law (37) expresses the protagonists' sanctions. When *turn = determiner* (the argumentation has ended), *winning = none* (the dispute is unresolved) and *winner = none* (no protagonist is declared winner), declaring *Protag* would be proper, timely and non-objectionable. If, however, *Protag* is sanctioned and *Protag'* is not, declaring *Protag* would be proper, timely but objectionable. (If both protagonists are sanctioned then objectionable declarations are defined only by laws (35), (36) and (38), that is, proper, timely declarations are non-objectionable.)

Similarly, we could have expressed the protagonists' sanctions in terms of objectionable claims, concessions, retractions, and so on. For example, we could have specified that the retractions of a sanctioned protagonist are objection-

able, even if proper and timely, and therefore, their effects could be blocked by objections.

In addition to dealing with the protagonists' forbidden actions, we want to reduce the possibility that the determiner will not comply with its obligation. Assuming that the determiner is obliged to declare *Protag* the winner of the dispute (see law (29)), the determiner may:

- declare *Protag'* the winner (more precisely, the last declaration before the last protocol timeout concerns *Protag'*),
- make no declaration, or
- declare *Protag* the winner.

In the third case the obligation is (temporarily) discharged. New evidence, however, may arise after the declaration (and before the last protocol timeout), such as an untimely retraction by *Protag*, obliging the determiner to declare *Protag'* the winner (by setting $winning = Protag'$) and this obligation may not be discharged.

Notice that the employed object mechanism is inadequate for ensuring a 'fair' result in the three aforementioned scenarios. (A 'fair' result implies that *Protag* is declared winner if the dispute is resolved in its favour.) In the first case, a well-founded objection to the action $declare(Det, Protag')$ will block the effects of the declaration but will not force $winner = Protag$. In the second case there is no declaration to object to, and in the third case the declaration was not objectionable at the time.

What sanctions should be enforced in these cases? There are a number of possibilities. We may specify, for instance, that a sanctioned determiner is disqualified from acting as a determiner in future argumentations, for a specified time period perhaps. If the argumentation takes place in the context of a computational system including several protocols, such as a system of negotiation or deliberative assemblies, a sanctioned determiner may have restricted permissions or its actions may be objectionable when participating in the remaining protocols, possibly occupying other roles. In addition, we may specify that it is proper and timely to initiate proceedings against a sanctioned determiner (assuming the existence of an adjudicating authority) in order to enforce compliance with the determiner's obligation, not as a 'punishment' to the determiner but as a way of discouraging it to avoid complying with the protocol rules and ensuring 'fairness'.

Sanctions are one means by which the performance of forbidden actions and non-compliance with obligations may be addressed. Another possible strategy is to devise physical controls that will force agents to comply with the protocol rules. For instance, repetitive objectionable actions may be physically blocked (since, in the present example, they are forbidden). The general strat-

egy of designing mechanisms to force compliance and eliminate non-permitted behaviour is what Jones and Sergot (1993) referred to as *regimentation*. Regimentation devices have often been employed in order to eliminate ‘anti-social’ behaviour in computational systems (see, for instance, (Klein et al., 2003; Minsky and Ungureanu, 2000; Rodriguez-Aguilar et al., 1998)). It has been argued (Jones and Sergot, 1993), however, that regimentation is rarely desirable (it results in a rigid system that may discourage agents from entering it (Prakken, 1998)), and not always practical. In any case, violations may still occur even when regimenting a computational system (consider, for instance, a faulty regimentation device). For all of these reasons, we have to allow for sanctioning and not rely exclusively on regimentation mechanisms.

7.5 Additional Considerations

We presented a formalisation of the RTFD* protocol rules. In order to perform computational experiments (such as those presented in Section 9), we need to code up (fragments of) the logic of disputation. The following law, for example, states that a protagonist accepts all (classical) logical implications of each of its premises:

$$\begin{aligned} \text{accepts}(\text{Protag}, Q) \text{ iff} \\ \text{premise}(\text{Protag}, P) = \mathbf{t} \wedge \text{implies}(P, Q) \end{aligned} \tag{39}$$

The *implies* here are simply suitably chosen rigid constants. In a similar manner, we may specify the acceptance of propositions as a result of the conjunctions of an agent’s premises and of premises regarding default rules.

We have formalised only a small fragment of the logic of disputation, enough to conduct simple experiments. The focus of this paper lies on the protocol rules rather than the logic of disputation (that, we assume, does not necessarily have to be prioritised default logic as in Brewka’s version). Moreover, a complete formalisation of the logic of disputation (on top of a formalisation of the protocol rules) would substantially increase the number of D^{RTFD^*} laws, thus significantly increasing CCALC’s time of computing answers to queries (regarding D^{RTFD^*}). These complications could be addressed by employing alternative action languages. We discuss this issue in Section 11.

Brewka, in his reconstruction of RTFD, places emphasis on the formalisation of the ‘silence implies consent’ principle: a protagonist that does not explicitly challenge a claim by the other protagonist is assumed to concede to the claim. We may incorporate this principle in RTFD* by modifying the formalisation

of the logic of disputation as follows:

$$\begin{aligned} \text{accepts}(\text{Protag}, Q) \text{ iff} \\ (\text{premise}(\text{Protag}, P) = \mathbf{t} \vee \text{premise}(\text{Protag}, P) = \mathbf{u}) \wedge \\ \text{implies}(P, Q) \end{aligned} \quad (40)$$

Recall that $\text{premise}(\text{Protag}, P) = \mathbf{u}$ expresses that *Protag* has an unconfirmed premise that P , that is, *Protag* has not responded to a claim that P made by the other protagonist. In a protocol adopting the ‘silence implies consent’ principle, a protagonist accepts all logical implications of each of its explicit *and unconfirmed* premises.

Note that if we do not want to incorporate the ‘silence implies consent’ principle then we disregard unconfirmed premises in the logic of disputation, by using law (39) instead of (40).

8 Proving Properties of RTFD*

The explicit transition systems semantics of the $\mathcal{C}+$ language enables us to prove various properties of the presented RTFD* specification (which is expressed by means of the definite action description D^{RTFD^*}). We may prove, for example, that as long as the determiner complies with its obligations the protocol result will be ‘fair’. Suppose we specify ‘fairness’ as follows:

$$\text{fair} \text{ iff } \text{winning} = \text{none} \vee \text{winner} = \text{winning} \quad (41)$$

The statically determined fluent constant *fair* holds in a protocol state if and only if the dispute is unresolved ($\text{winning} = \text{none}$) or the dispute is resolved in favour of the declared winner ($\text{winner} = \text{winning}$). We want to examine whether or not a protocol result is ‘fair’, that is, whether or not the fluent constant *fair* is true in the final state of a protocol execution.

Proposition 1 *The protocol result will be ‘fair’ if and only if in the final protocol state there is no obligation on the determiner to declare.*

PROOF. First we will prove that if in the final protocol state there is no obligation on the determiner then the protocol result will be ‘fair’. Assume a final protocol state s in which there is no obligation on the determiner and the result is not ‘fair’:

$$\begin{aligned} s \not\models \text{obl}(\text{declare}(\text{det}, \text{Protag})), \text{ for any } \text{Protag} \\ s \models \text{winning} \neq \text{none} \wedge \text{winner} \neq \text{winning} \wedge \text{turn} = \text{role_of}(\text{det}) \end{aligned}$$

(Recall that $role_of(det) = determiner$ is ‘rigid’. Moreover, in a final protocol state $turn = determiner$.) $s \models winning \neq none$ iff $s \models winning = Protag$ for some protagonist $Protag$, and so:

$$s \models winning = Protag \wedge winner \neq Protag$$

Since s is a state of D^{RTFD^*} , it is an interpretation of σ^f such that

$$s = T_{static}(s) \cup Simple(s)$$

$T_{static}(s) =_{\text{def}} \{F \mid \text{static law ‘}F \text{ if } G\text{’ is in } D^{RTFD^*}, s \models G\}$ and $Simple(s)$ denotes the set of simple fluent atoms satisfied by s (see Section 2.3). From law (29) and the fact that

$$s \models turn = role_of(det) \wedge winning = Protag \wedge winner \neq Protag$$

we have that

$$obl(declare(det, Protag)) \in T_{static}(s)$$

According to our initial assumption, however, there is no obligation on the determiner in s , which implies that $s \neq T_{static}(s) \cup Simple(s)$. Therefore, s is not a state of D^{RTFD^*} .

Second we will prove that if the protocol result is ‘fair’ then in the final state there is no obligation on the determiner to declare. Assume a final protocol state s in which there is an obligation on the determiner to declare and the result is ‘fair’. If $fair$ is true in s then either $winning = none$ is in s or $winner = winning$ is in s . Consider first the case $winning = none$:

$$s \models obl(declare(det, Protag)) \wedge winning = none \wedge turn = role_of(det)$$

for some $Protag$. Given law (29), we have that

$$obl(declare(det, Protag)) \notin T_{static}(s)$$

Moreover, since obl are statically determined fluent constants, we also have that

$$obl(declare(det, Protag)) \notin Simple(s)$$

Therefore, $s \neq T_{static}(s) \cup Simple(s)$ and s is not a state of D^{RTFD^*} .

The proof that s is not a state of D^{RTFD^*} for the case in which $winner = winning$ is similar. ■

The obligation to declare might never arise during a protocol execution. Proposition 1 shows that in this case the result will be ‘fair’. If an obligation to

declare does arise, however, the protocol result is also guaranteed to be ‘fair’, as long as the determiner complies with the obligation.

CCALC provides an automated means for proving properties of a protocol specification. We express the $\mathcal{C}+$ action description D^{RTFD^*} in CCALC’s input language and then query CCALC about D^{RTFD^*} to prove properties of the RTFD* specification. (The types of query that CCALC computes were presented in Section 3. Details of CCALC’s input language may be found in (Lee et al., 2001; Akman et al., 2004).) Consider the following example:

Proposition 2 *A proper, timely concede(Protag, Q) action always leads to a state in which Protag has an explicit premise that Q.*

We instruct CCALC to compute all states s' such that

- (s, ε, s') is a transition of D^{RTFD^*} ,
- $s \models \text{proper}(\text{concede}(\text{Protag}, Q)) \wedge \text{timely}(\text{concede}(\text{Protag}, Q))$, and
- $\varepsilon \models \text{concede}(\text{Protag}, Q)$.

For every state s' computed by CCALC we obtain

$$s' \models \text{premise}(\text{Protag}, Q) = \mathbf{t}$$

This is because, briefly, a proper and timely concession is not objectionable in the presented RTFD* (see Table 3). Therefore, the effects of such an action, expressed by law (15), cannot be blocked. In other words, even when

$$\varepsilon \models \text{concede}(\text{Protag}, Q) \wedge \text{objected}(Ag)$$

the resulting state s' always includes $\text{premise}(\text{Protag}, Q) = \mathbf{t}$.

More details about the computational experiments performed with CCALC on D^{RTFD^*} are presented in the following section.

We may prove further properties of the RTFD* specification, in the manner shown above, such as that the effects of any non-objectionable action cannot be blocked, an objection always blocks the effects of an objectionable action, there is always a permitted action for a participant when it is its turn to ‘speak’, a protagonist’s performance of a forbidden action always leads to a sanction, the determiner is never obliged and forbidden to declare the winner, and so on.

In Section 11 we discuss alternative techniques for proving properties of a protocol specification.

9 Executing RTFD*

Proving properties of a specification can be seen as a design-time activity. For instance, protocol designers may wish to prove properties of a protocol specification in order to determine whether or not this specification meets their requirements. Additionally, agents (or their designers) may wish to prove various properties of a protocol specification when deciding whether to enter (deploy their agents in) that protocol.

At run-time, we may execute a protocol specification to provide, amongst other things, information about the protocol state current at each time. Computation of such information is a special case of a prediction query (see Section 3). A protocol state — which actions are proper, timely, objectionable, permitted, and so on — may be publicised to (a subset of) the protocol participants, or their designers. (Such run-time services may be provided by a central server or in various distributed configurations. Further discussion of these architectural issues is outside the scope of this paper.) Other run-time services include the calculation of plans, by means of computing answers to planning queries, and retrieval of past protocol states, by means of computing answers to postdiction queries (as already mentioned in Section 3, plans may be additionally computed at design-time). We show in this section example prediction and planning queries on the RTFD* specification, and the results obtained. To save space, details of postdiction query examples are omitted.

To execute the presented RTFD* protocol, one has to choose specific values for the following parameters:

- Duration of timeouts; this can be expressed as the maximum number of messages that can be exchanged in the given communication channel in the interval defined by any two consecutive timeout events. In our formalisation a timeout duration is expressed as the maximum number of transitions that may take place between any two consecutive timeout events (in this specification a transition is labelled with either a main protocol action, objected or not, or a timeout event).
- Number of turns for each protagonist.
- Number of permitted objectionable actions (see Section 7.3).

To conduct computational experiments, arbitrary numerical values were chosen for these parameters. For a concrete illustration we will present here experiments in which at most two main protocol actions are physically possible between two consecutive timeout events (that is, at most two transitions may take place between two consecutive timeout events), each protagonist has three turns to ‘speak’, and one objectionable action is permitted. As already mentioned, the specified logic of disputation is very simple (only classical im-

Table 5
A Sample Run of RTFD*

act	proper	timely	obj	per	obl
<i>cl(pro, q)</i>	<i>cl(Protag, p)</i> <i>cn(opp, q)</i> <i>rtr(pro, q)</i> <i>dn(opp, q)</i> <i>dcl(det, pro)</i>	<i>cl(pro, Q)</i> <i>cn(pro, Q)</i> <i>rtr(pro, Q)</i> <i>dn(pro, Q)</i>	<i>cl(opp, p)</i> <i>cn(opp, q)</i> <i>dn(opp, q)</i> <i>dcl(det, Protag)</i>	all	none
<i>timeout</i>	<i>cl(Protag, p)</i> <i>cn(opp, q)</i> <i>rtr(pro, q)</i> <i>dn(opp, q)</i> <i>dcl(det, pro)</i>	<i>cl(opp, Q)</i> <i>cn(opp, Q)</i> <i>rtr(opp, Q)</i> <i>dn(opp, Q)</i>	<i>cl(pro, p)</i> <i>rtr(pro, q)</i> <i>dcl(det, Protag)</i>	all	none
<i>dn(opp, q), objected(pro)</i>	<i>cl(pro, p)</i> <i>cl(opp, Q)</i> <i>rtr(pro, q)</i> <i>dcl(det, Protag)</i>	<i>cl(opp, Q)</i> <i>cn(opp, Q)</i> <i>cn(opp, Q)</i> <i>dn(opp, Q)</i>	<i>cl(pro, p)</i> <i>rtr(pro, q)</i> <i>dcl(det, Protag)</i>	all	none
<i>cl(pro, p), objected(opp)</i>	<i>cl(pro, p)</i> <i>cl(opp, Q)</i> <i>rtr(pro, q)</i> <i>dcl(det, Protag)</i>	<i>cl(opp, Q)</i> <i>cn(opp, Q)</i> <i>cn(opp, Q)</i> <i>dn(opp, Q)</i>	<i>cl(pro, p)</i> <i>rtr(pro, q)</i> <i>dcl(det, Protag)</i>	<i>cl(pro, q)</i> <i>cl(opp, Q)</i> <i>cn(Protag, Q)</i> <i>rtr(pro, p)</i> <i>rtr(opp, Q)</i> <i>dn(Protag, Q)</i> <i>dcl(det, Protag)</i>	none

plication is allowed). Moreover, for the experiments presented ‘silence implies consent’ *is* incorporated in the protocol specification.

We mentioned earlier that the propositions that the two protagonists may claim, concede to, retract, and deny should also be specified at the outset. (As will be discussed later, this restriction may be lifted by re-compiling D^{RTFD^*} each time a protagonist claims a new proposition.) In order to keep sample runs small we will present queries and the computed results concerning only two propositions: the topic q , and p .

The first column of Table 5 shows a sample run of RTFD*; the information displayed on the remaining columns, that is, which actions are proper, timely, objectionable, permitted and obligatory at each state, is produced by computing answers to a number of prediction queries. To save space, in Table 5 *claim* is written as *cl*, *concede* as *cn*, *retract* as *rtr*, *deny* as *dn* and *declare* as *dcl*. Terms containing variables *Protag* and *Q* stand for all their instances.

Consider the following prediction query: at the initial protocol state the proponent *pro* claims the topic of the argumentation *q*; which actions are proper, timely, objectionable, permitted or obligatory at the resulting state? The answer to this query is displayed in the six rows below the $claim(pro, q)$ action labelled $cl(pro, q)$.

At the state resulting from $claim(pro, q)$, proponent *pro* has a premise that *q* (see law (13)). Therefore, at that state, it is improper for *pro* to claim *q* (see law (20)). It is proper for *pro*, however, to claim *p* because, at that state, *pro* does not have a premise that *p*. The remaining information displayed in Table 5 is computed in a similar manner.

Note that at the state resulting from $claim(pro, q)$ it is timely only for *pro* to perform an action. This is so because it is the proponent's turn to speak (see Table 2 for the specification of timely actions).

$claim(opp, p)$, $concede(opp, q)$ and $deny(opp, q)$ are objectionable in the state reached by the performance of $claim(pro, q)$. This is due to the fact that these actions are untimely and proper. At the same state declaring either *pro* or *opp* is objectionable because in either case the declaration is untimely and $winner = none$ (see Table 3 for the specification of objectionable actions).

All actions are permitted at the resulting state of the query presented above. Recall that an action is permitted if it is not objectionable or the specified limit of objectionable actions has not been reached; in this sample run one objectionable action per participant is permitted. In the same state, no action is obligatory; obligations are associated with declarations and arise when it is the determiner's turn to speak. (The specifications of permissions and obligations were presented in Table 4.)

The next action in the protocol run presented in Table 5 is a timeout that sets $turn = opponent$. Consequently, in the following state it is timely only for *opp* to perform an action. After the timeout *opp* denies *q* while *pro* objects to *opp*'s action. The denial is not objectionable and, therefore, *pro*'s objection does not block the effects of *opp*'s action. The next action in the protocol run, however, $claim(pro, p)$, is objectionable and thus *opp*'s objection blocks the effects of *pro*'s action.

After the performance of *pro*'s objectionable claim, *pro* is no longer permitted to perform an objectionable action because it reached the specified limit of permitted objectionable actions. Non-conformance with this prohibition will sanction *pro*.

A final remark on the protocol run presented in Table 5 concerns the 'silence implies consent' principle. Due to this principle, at the first two states displayed in Table 5 it is proper to declare *pro*, but improper to declare *opp*, the

Table 6
A Sample Run of RTFD*

act	proper	timely	obj	per	obl
<i>timeout</i>	$cl(Protag, Q)$ $dcl(det, opp)$	$dcl(det, Protag)$	$cl(Protag, Q)$ $dcl(det, pro)$	$cl(Protag, Q)$ $dcl(det, opp)$	$dcl(det, opp)$
$dcl(det, pro)$	$cl(Protag, Q)$ $dcl(det, opp)$	$dcl(det, Protag)$	$cl(Protag, Q)$	$cl(Protag, Q)$ $dcl(det, opp)$	$dcl(det, opp)$
$dcl(det, opp)$, $objected(pro)$	$cl(Protag, Q)$	$dcl(det, Protag)$	$cl(Protag, Q)$ $dcl(det, pro)$	$cl(Protag, Q)$ $dcl(det, Protag)$	none

argumentation winner. In these states *pro* is ‘winning’ the dispute because $premise(pro, q) = t$ and $premise(opp, q) = u$ which imply, due to ‘silence implies consent’ (see law (40)), that both protagonists accept *q*, the topic of argumentation (the *winning* fluent constant is defined by laws (22)–(24)). If ‘silence implies consent’ were not incorporated in the protocol then $premise(opp, q) = u$ would not imply that *opp* accepts the topic (see law (39)), and thus no protagonist would be ‘winning’. In this case, it would be proper to declare either *pro* or *opp* the argumentation winner.

Table 6 presents another sample protocol run of RTFD* and the associated information produced by computations of query answers. The timeout sets $turn = determiner$. The proponent *pro* does not accept the argumentation topic and thus the determiner is obliged to declare the opponent *opp* the winner. Clearly, the next action of the presented run, $declare(det, pro)$, does not discharge this obligation. The obligation is discharged when $declare(det, opp)$, although *pro* objects to this declaration; *pro*’s objection does not block the effects of the declaration as the latter action is not objectionable.

In addition to producing the protocol state current at each time, we may compute plans in order to facilitate agents to achieve their goals. Consider the following planning query: we are in a state in which: (i) it is the determiner’s turn to speak, (ii) the dispute is unresolved, that is, no protagonist is ‘winning’, and (iii) there is no declared winner. Find all paths to a final protocol state, that is, a state reached after the final timeout, in which the proponent is declared the winner.

CCALC finds several solutions to this query. All solutions include the action $declare(det, pro)$ and the final timeout. Moreover, the declaration is either not objected, or the proponent is not the only protagonist sanctioned. If the proponent was sanctioned and the opponent was not, the $declare(det, pro)$ action would be objectionable (see law (37)) and thus an objection would block its effects.

Here is another planning query: given the initial state of the protocol, is it possible to reach a state, within the maximum number of transitions of D^{RTFD^*} , in which an agent: (i) has exceeded the limit of permitted objectionable actions and, (ii) is not sanctioned?

CCALC finds no solution within the maximum number of transitions of D^{RTFD^*} . Given the chosen values for the number of turns of each protagonist (three), and the maximum number of transitions that may take place between any two consecutive timeouts (two), it can be calculated that the maximum number of transitions of D^{RTFD^*} , that is, the number of transitions of the longest path of D^{RTFD^*} , is twenty-one. (Intuitively, in the longest path of D^{RTFD^*} the proponent performs two actions and then a timeout takes place, signalling the opponent's turn; the opponent then performs two actions, followed by a timeout signalling the proponent's turn. The aforementioned sequence of actions is repeated three times, that is, the number of turns of each protagonist. Finally, the determiner performs two actions and the last timeout takes place.) Since there is no solution within the maximum number of transitions, starting from the initial protocol state, we will never reach a state in which an agent has exceeded the limit of permitted objectionable actions and is not sanctioned.

10 Related Work

In Section 6 we discussed the points of departure of our formalisation from Brewka's account. Briefly:

- We employed the $\mathcal{C}+$ language to specify the argumentation protocol, instead of the Situation Calculus. Moreover, we *executed* the protocol specification with the use of CCALC, thus providing several design-time and run-time services to protocol designers, agent designers and agents themselves.
- We introduced deadlines to cater for realistic multi-agent protocols.
- We refined Brewka's distinction of possible and legal actions. In our formalisation an action can be classified as physically possible, proper, timely, objectionable, permitted or obligatory.
- We maintained the object mechanism; however, we forbid multiple objectionable actions to avoid having to object to each such action.
- We introduced sanctions in order to discourage participants from performing forbidden actions and not complying with obligations.
- We retained the 'silence implies consent' principle, as an *optional* feature of the protocol.

Apart from 'static' argument systems, Brewka (2001, Section 6) formalises 'dynamic' argument systems, that is, argument systems in which participants

can start a meta-level debate, arguing about the protocol rules. ‘Dynamic’ protocol specifications are out of the scope of this paper.

Argumentation protocols have long been studied in the fields of philosophy and computer science — see (Prakken, 2006) for a recent review. The focus of this paper was on the procedural part of argumentation, that is, we focused on the specification of the protocol rules rather than the logic of disputation. This is in contrast to research in nonmonotonic and uncertain reasoning argument systems that, as Brewka (2001) mentions, have been used to define inference systems for existing nonmonotonic logics, or a non-standard consequence relation for logics based on a notion of argument. (See, for example, (Cayrol, 1995; Dung, 1995; Krause et al., 1995; Pollock, 1996; Bondarenko et al., 1997; Prakken and Sartor, 1997); Chesñevar et al. (2000), and Prakken and Vreeswijk (2002) provide two surveys.) Different logics of disputation are suited to different types of argument. Our protocol formalisation was not bound to a specific logic of disputation.

Like (Gordon, 1994, 1995; Prakken, 2000, 2005; Brewka, 2001), we adopted a ‘public semantics’ (Prakken, 2006, Section 6), that is, we made no assumptions about the participants’ internal architectures. This is contrast to approaches that allow for protocol rules referring to a participant’s internal belief base (for instance, (Kraus et al., 1998; Amgoud et al., 2000; Parsons and McBurney, 2003; Parsons et al., 2002, 2003a,b)).

A line of research that is closely related to our work is that of (Bodenstaff et al., 2006). These researchers employ Shanahan’s (1999) ‘full Event Calculus’ to formalise Prakken’s (2005) dialogue system for argumentation, and Parsons, Wooldridge and Amgoud’s (2003b) persuasion dialogue. Bodenstaff et al. formalise the procedural aspect of argumentation, expressing the ‘legal’ protocol moves in terms of a ‘reply structure’, distinguishing between ‘attacking’ and ‘surrendering’ replies. In this paper we presented a finer classification of protocol actions than Bodenstaff and colleagues’ classification of ‘legal’/‘illegal’ actions. Concerning the reply structure, Prakken (2006) notes that it is not a standard feature of all dialogue systems (see, for example (Fulda, 2001)). In any case, it is possible to adjust our formalisation in order to express a reply structure (in this paper our aim was to reconstruct Brewka’s account of RTFD that did not explicitly include such a structure). This can be done by adjusting our specification of proper actions.

Apart from research on argumentation, work that has similar objectives to ours comes from the distributed artificial intelligence literature on norm-governed systems specification. A few examples are the approaches on ‘artificial social systems’ (Moses and Tennenholtz, 1992, 1995; Shoham and Tennenholtz, 1992, 1995; Tennenholtz, 1995; Fitoussi and Tennenholtz, 2000), ‘law-governed interaction’ (Minsky and Ungureanu, 2000), and ‘electronic institutions’ (Es-

teva et al., 2000, 2001, 2002b,a; Rodriguez-Aguilar and Sierra, 2002). Close to our work is Yolum and Singh’s (2002; 2004) work on ‘commitment protocols’. These researchers formalise, in Shanahan’s ‘full Event Calculus’, a set of operations on commitments such as create, discharge, cancel, release, and so on. Moreover, they employ an Event Calculus planner (Shanahan, 2000) to facilitate the planning of commitment protocol participants. It is important to note that in Yolum and Singh’s work the term ‘commitment’ refers to a form of (directed) obligation between agents, and is *not* used as an alternative term for ‘premise’. It is difficult to see how an argumentation protocol, or many other interaction protocols for multi-agent systems (for instance, protocols for negotiation, voting, performing transactions in electronic marketplaces, and so on), can be specified simply in terms of commitments in this sense. At the very least, a specification of a protocol’s constitutive norms is also required.

11 Conclusion

We have focussed in this paper on the formal representation of the procedural aspects of an argumentation protocol, using Brewka’s reconstruction of Rescher’s theory of formal disputation as a concrete example. We distinguished in the specification between the constitutive rules defining the protocol itself — the protocol actions and their effects on the protocol state — and the physical environment within which the protocol is executed, and the normative environment which may place further constraints on what protocol actions are permitted or obligatory. Although in the simplest cases a protocol action is permitted if and only if it is both proper and timely, there are many reasons why this need not always be so. We presented some simple examples by way of illustration, but it should be clear that there are many other possibilities that we did not discuss. There is also much choice in deciding where to place the boundary between the constitutive and normative components. For instance, we chose in the example presented here to say that the claim of Q by a protagonist who already holds Q as a premise is not proper, has no effect on the protocol state, but is permitted. We could have chosen to say instead that such a claim is proper but never permitted. Or that it is not proper, does have an effect on the protocol state (in that it requires the other protagonist to concede or deny), but is not permitted. One objective of the present work is to allow such variations to be specified and examined, and to help evaluate the effectiveness of proposed sanctioning and enforcement mechanisms.

In an earlier formalisation of Brewka’s protocol (Artikis et al., 2003) we specified the constitutive elements of the protocol by defining the conditions under which an agent has institutional power (competence, capacity) to perform a particular protocol action; we then said that an action is ‘valid’ when it is performed by an empowered agent. In this paper we have constructed a

more structured and detailed specification by defining separately the conditions under which an action is proper and timely, as has also been suggested by Prakken et al. (Prakken, 1998; Prakken and Gordon, 1999). One aim of the paper was to see how this additional structure would be reflected in the specification of a concrete example. A more detailed specification still would define what exact forms of message or utterance count as expressing a claim, concede, deny, retract, declare, and object action. We have not defined that level of detail here. It would be an important component in a run-time mechanism but is not so important if we are primarily interested, as here, in investigating properties of the specified protocol. We have also found that it adds flexibility to specify separately which actions are objectionable, though that is perhaps difficult to demonstrate convincingly in a comparatively simple protocol such as Brewka's. Finally, we have shown how the 'silence implies consent' principle can be included straightforwardly as an optional component by marking as unconfirmed those premises that have not been explicitly conceded or retracted by the opponent.

In principle, the kind of specifications presented in this paper could also be expressed in other temporal reasoning formalisms. The $\mathcal{C}+$ language, however, has a number of important features that have led us to choose it as the basis for further developments. First, it is a comparatively expressive formalism with fine control for specifying default persistence of fluent constants. The availability of static laws moreover means, amongst other things, that complex specifications can be given structure. For example, most of the rules defining *proper(...)* and *timely(...)* constants in this paper can be expressed as static laws with simple conjunctions as their conditions. Additional structure can be provided by introducing suitably chosen intermediate concepts, such as the fluent *winning* used in our example, themselves defined by means of static laws. This ability is important if large specifications are to be undertaken.

Second, besides its semantics through translation to the formalism of 'non-monotonic causal theories' (Giunchiglia et al., 2004), a $\mathcal{C}+$ action description has an explicit semantics in terms of transition systems. This is important because it provides a link to a wide range of other formalisms and tools based on transition systems. We have been able to devise, for example, an extended form of $\mathcal{C}+$ specifically designed for representing norms and institutions (Sergot, 2004a,b; Sergot and Craven, 2006), including direct support for (a version of) the 'counts as' relation for actions (Jones and Sergot, 1996) and a treatment of permitted/forbidden states, transitions and runs. The relationship between permitted and obligatory actions presented in this paper, for instance, which had to be formulated explicitly in the $\mathcal{C}+$ specification (see Section 7.3), is built into the semantics of the language in the extended version. We have not employed the extended form of $\mathcal{C}+$ in this paper because space prevents the presentation of its additional features.

The language $\mathcal{C}+$ (and its derivatives) also has some important limitations. Most obviously, from a representational point of view, the language inherits the limitations of transition systems, in particular that the executable actions (transitions) in any given state s of the system, and their effects, can depend only on the state s and not on the path or history by which state s was reached (unless of course we encode the entire history in every state s). The $\mathcal{C}+$ language itself (though not the underlying formalism of causal theories) can only express causes relationships between successive states; delayed effects cannot be expressed directly. This is what makes a detailed formulation of the object mechanism awkward and fiddly to express in $\mathcal{C}+$, and why we chose to omit the details from this paper. (But see (Craven and Sergot, 2005) for a modified form of $\mathcal{C}+$ which can express delayed effects.) In other protocols where the moves available to a participant might depend on the entire history of the protocol so far (a participant is restricted on the number of repeated claims it can make, for instance), these limitations are not so easily overcome. Extensions to $\mathcal{C}+$ capable of expressing such constraints in a concise manner are one direction of our current research.

From the point of view of implementation, CCALC provides an immediate and convenient means of implementing $\mathcal{C}+$ action descriptions. The execution of the RTFD* specification, however, confirmed our previous experience regarding CCALC's efficiency, and in particular that it does not provide a practical means for supporting run-time activities. A major limitation, of course, is the need to encode the underlying logic of disputation, which here we have done by means of explicit instances of the *implies* fluent constants. This clearly works only for small examples. It also requires that we are able to specify at the outset the complete set of propositions and default rules that could be claimed during the argumentation. Whilst it is possible to improve the efficiency of the CCALC implementation by pre-compiling this part of the action description (the *implies* constants are 'rigid'), and even in principle re-compiling the action description every time a protagonist claims a new, unanticipated, proposition, this is clearly not a practical way of supporting run-time activities. It is only adequate for conducting computational experiments with the specification of the kind discussed in Section 9. (The 'logic of disputation' for the examples in that section is trivial but that was chosen deliberately to keep the sample runs small enough to be presented.)

CCALC is not the only means by which $\mathcal{C}+$ action descriptions could be executed. We have also used versions of the Event Calculus to specify and execute (an earlier version of) the RTFD* protocol (Artikis, 2003, Sections 6.10–6.12). Given an instance of the specification and a narrative — a record of what actions have been performed so far — this (Prolog) implementation allows all protocol states, including what is permitted and obligatory at each state, to be queried and computed efficiently. It is not necessary to specify all propositions in advance, and it is comparatively easy to implement the underlying

logic of disputation as another Prolog module to be added to the protocol specification. The Event Calculus implementation, however, is not well suited to planning and postdiction tasks. More importantly, we also lose the explicit transition system semantics which we see as the single most important advantage of the $\mathcal{C}+$ formulation.

Recent work by Craven (2006) has investigated the relationships between Event Calculus and $\mathcal{C}+$. His $\mathcal{EC}+$ implementation provides an efficient Event Calculus style of computation of narratives with (a restricted form of) the $\mathcal{C}+$ language, providing a promising means of supporting run-time activities. He has also exploited the transition system semantics to connect $\mathcal{C}+$ to model checking software (specifically NuSMV (Cimatti et al., 2002)). This allows protocol properties, expressed in Linear Temporal Logic (LTL) and Computation Tree Logic (CTL), to be verified by means of standard model checking techniques on protocol specifications expressed in the $\mathcal{C}+$ language. Application of these methods to larger examples is a direction of our current research.

Acknowledgements

This article is an updated and extended version of (Artikis et al., 2003). We are grateful to the reviewers and the participants of ICAIL'03, who gave us useful feedback. We should also like to thank Joohyung Lee and Vladimir Lifschitz for their suggestions regarding the $\mathcal{C}+$ language and CCALC. The authors themselves, however, are solely responsible for any misunderstanding about the use of these technologies.

This work was supported partly by the EU ALFEBIITE Project (IST-1999-10298) and partly by the EPSRC project 'Theory and Technology of Norm-Governed Self-Organising Networks' (GR S74911 01).

References

- Akman, V., Erdogan, S., Lee, J., Lifschitz, V., Turner, H., 2004. Representing the zoo world and the traffic world in the language of the Causal Calculator. *Artificial Intelligence* 153 (1–2), 105–140.
- Amgoud, L., Maudet, N., Parsons, S., 2000. Modelling dialogues using argumentation. In: *Proceedings of the International Conference on Multiagent Systems (ICMAS)*. IEEE Computer Society, pp. 31–38.
- Artikis, A., November 2003. Executable specification of open norm-governed computational systems. Ph.D. thesis, University of London, retrieved March 6, 2006, from <http://www.doc.ic.ac.uk/~aartigis/publications/artikis-phd.pdf>, also available from the author.

- Artikis, A., Sergot, M., Pitt, J., 2003. An executable specification of an argumentation protocol. In: Proceedings of Conference on Artificial Intelligence and Law (ICAIL). ACM Press, pp. 1–11.
- Artikis, A., Sergot, M., Pitt, J., 2006. Specifying norm-governed computational societies. Tech. Rep. 2006/5, Imperial College London, Department of Computing, retrieved March 6, 2006, from <http://www.doc.ic.ac.uk/research/technicalreports/2006/DTR06-5.pdf>.
- Bodenstaff, L., Prakken, H., Vreeswijk, G., 2006. On formalising dialogue systems for argumentation in the event calculus. In: Proceedings of Workshop on Non-Monotonic Reasoning.
- Bondarenko, A., Dung, P. M., Kowalski, R., Toni, F., 1997. An abstract, argumentation-theoretic approach to default reasoning. *Artificial Intelligence* 93, 63–101.
- Brewka, G., 2001. Dynamic argument systems: a formal model of argumentation processes based on situation calculus. *Journal of Logic and Computation* 11 (2), 257–282.
- Brewka, G., Eiter, T., 1998. Prioritizing default logic. In: Festschrift 60th Anniversary of W. Bibel. Kluwer.
- Cayrol, C., 1995. On the relation between argumentation and coherence based entailment. In: Proceedings of International Joint Conference on Artificial Intelligence (IJCAI). pp. 1443–1448.
- Chesñevar, C., Maguitman, A., Loui, R., 2000. Logical models of argument. *ACM Computing Surveys* 32 (4), 337–383.
- Cimatti, A., Clarke, E., Giunchiglia, E., Giunchiglia, F., Pistore, M., Roveri, M., Sebastiani, R., Tacchella, A., 2002. NuSMV Version 2: An OpenSource Tool for Symbolic Model Checking. In: Proc. International Conference on Computer-Aided Verification (CAV 2002), Copenhagen, July 2002. LNCS 2404. Springer, see <http://nusmv.irst.itc.it>.
- Craven, R., September 2006. Execution mechanisms for the action language $\mathcal{C}+$. Ph.D. thesis, University of London.
- Craven, R., Sergot, M., 2005. Distant causation in $\mathcal{C}+$. *Studia Logica* 79 (1), 73–96.
- Dung, P. M., 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence* 77 (2), 321–358.
- Esteva, M., de la Cruz, D., Sierra, C., 2002a. ISLANDER: an electronic institutions editor. In: Castelfranchi, C., Johnson, L. (Eds.), Proceedings of Conference on Autonomous Agents and Multi-Agent Systems (AAMAS). ACM Press, pp. 1045–1052.
- Esteva, M., Padget, J., Sierra, C., 2002b. Formalizing a language for institutions and norms. In: Meyer, J.-J., Tambe, M. (Eds.), *Intelligent Agents VIII: Agent Theories, Architectures, and Languages*. LNAI 2333. Springer, pp. 348–366.
- Esteva, M., Rodriguez-Aguilar, J., Arcos, J., Sierra, C., Garcia, P., 2000. Institutionalising open multi-agent systems. In: Durfee, E. (Ed.), Proceedings

- of the International Conference on Multi-agent Systems (ICMAS). IEEE Press, pp. 381–382.
- Esteva, M., Rodriguez-Aguilar, J., Sierra, C., Garcia, P., Arcos, J., 2001. On the formal specifications of electronic institutions. In: Dignum, F., Sierra, C. (Eds.), *Agent Mediated Electronic Commerce*. LNAI 1991. Springer, pp. 126–147.
- Fitoussi, D., Tennenholtz, M., 2000. Choosing social laws for multi-agent systems: minimality and simplicity. *Artificial Intelligence* 119 (1-2), 61–101.
- Fulda, J., 2001. The logic of ‘improper cross’. *Artificial Intelligence and Law* 8 (4), 337–341.
- Giunchiglia, E., Lee, J., Lifschitz, V., McCain, N., Turner, H., 2004. Nonmonotonic causal theories. *Artificial Intelligence* 153 (1–2), 49–104.
- Giunchiglia, E., Lee, J., Lifschitz, V., Turner, H., 2001. Causal laws and multi-valued fluents. In: *Proceedings of Workshop on Nonmonotonic Reasoning, Action and Change (NRAC)*.
- Gordon, T., 1994. The pleadings game: an exercise in computational dialectics. *Artificial Intelligence and Law* 2, 239–292.
- Gordon, T., 1995. *The Pleadings Game: An Artificial Intelligence Model of Procedural Justice*. Kluwer Academic Publishers.
- Jones, A., Sergot, M., 1993. On the characterisation of law and computer systems: the normative systems perspective. In: *Deontic Logic in Computer Science: Normative System Specification*. J. Wiley and Sons, pp. 275–307.
- Jones, A., Sergot, M., 1996. A formal characterisation of institutionalised power. *Journal of the IGPL* 4 (3), 429–445.
- Klein, M., Rodriguez-Aguilar, J., Dellarocas, C., 2003. Using domain-independent exception handling services to enable robust open multi-agent systems: the case of agent death. *Journal of Autonomous Agents and Multi-Agent Systems* 7 (1–2), 179–189.
- Kowalski, R., Sergot, M., 1986. A logic-based calculus of events. *New Generation Computing* 4 (1), 67–96.
- Kraus, S., Sycara, K., Evenchik, A., 1998. Reaching agreements through argumentation: a logical model and implementation. *Artificial Intelligence* 104 (1-2), 1–69.
- Krause, P., Ambler, S., Fox, J., 1995. A logic of argumentation for uncertain reasoning. *Computational Intelligence* 11 (1), 113–131.
- Lee, J., Lifschitz, V., Turner, H., 2001. A representation of the zoo world in the language of the Causal Calculator. In: *Proceedings of Symposium on Formalizations of Commonsense Knowledge*.
- Lifschitz, V., 2000. Missionaries and cannibals in the Causal Calculator. In: Cohn, A., Giunchiglia, F., Selman, B. (Eds.), *Proceedings of Conference on Principles of Knowledge Representation and Reasoning (KR)*. Morgan Kaufmann, pp. 85–96.
- Lifschitz, V., McCain, N., Remolina, E., Tacchella, A., 2000. Getting to the airport: the oldest planning problem in AI. In: Minker, J. (Ed.), *Logic-Based Artificial Intelligence*. Kluwer, pp. 147–168.

- Makinson, D., 1986. On the formal representation of rights relations. *Journal of Philosophical Logic* 15, 403–425.
- Minsky, N., Ungureanu, V., 2000. Law-governed interaction: a coordination and control mechanism for heterogeneous distributed systems. *ACM Transactions on Software Engineering and Methodology (TOSEM)* 9 (3), 273–305.
- Moses, Y., Tennenholtz, M., 1992. On computational aspects of artificial social systems. In: *Proceedings of Workshop on Distributed Artificial Intelligence (DAI)*. pp. 267–284.
- Moses, Y., Tennenholtz, M., 1995. Artificial social systems. *Computers and Artificial Intelligence* 14 (6), 533–562.
- Parsons, S., McBurney, P., 2003. Argumentation-based communication between agents. In: *Communication in Multiagent Systems*. LNCS 2650. Springer, pp. 164–178.
- Parsons, S., Wooldridge, M., Amgoud, L., 2002. An analysis of formal inter-agent dialogues. In: *Proceedings of Conference on Autonomous Agents and Multi-Agent Systems*. ACM Press, pp. 394 – 401.
- Parsons, S., Wooldridge, M., Amgoud, L., 2003a. On the outcomes of formal inter-agent dialogues. In: *Proceedings of Conference on Autonomous Agents and Multi-Agent Systems*. ACM Press, pp. 616–623.
- Parsons, S., Wooldridge, M., Amgoud, L., 2003b. Properties and complexity of some formal inter-agent dialogues. *Journal of Logic and Computation* 13 (3), 347–376.
- Pollock, J., 1996. Oscar - a general purpose defeasible reasoner. *Journal of applied non-classical logics* 6 (1), 89–113.
- Prakken, H., 1998. Formalising Robert’s rules of order. Tech. Rep. 12, GMD – German National Research Center for Information Technology.
- Prakken, H., 2000. On dialogue systems with speech acts, arguments, and counterarguments. In: *Proceedings of Workshop on Logics in Artificial Intelligence*. LNAI 1919. Springer, pp. 224–238.
- Prakken, H., 2005. Coherence and flexibility in dialogue games for argumentation. *Journal of Logic and Computation* 15, 1009–1040.
- Prakken, H., 2006. Formal systems for persuasion dialogue. *Knowledge Engineering Review* 21 (2), 163–188.
- Prakken, H., Gordon, T., 1999. Rules of order for electronic group decision making – a formalization methodology. In: Padget, J. (Ed.), *Collaboration between Human and Artificial Societies*. LNCS 1624. Springer, pp. 246–263.
- Prakken, H., Sartor, G., 1997. Argument-based extended logic programming with defeasible priorities. *Journal of Applied Non-Classical Logics* 7 (1), 25–75.
- Prakken, H., Vreeswijk, G., 2002. Logics for defeasible argumentation. In: Gabbay, D., Guenther, F. (Eds.), *Handbook of philosophical logic*. Vol. 4. Kluwer Academic Publishers, pp. 218–319.
- Reiter, R., 2001. *Knowledge in Action: Logical Foundations for Describing and Implementing Dynamical Systems*. The MIT Press.
- Rescher, N., 1977. *Dialectics: A Controversy-Oriented Approach to the Theory*

- of Knowledge. State University of New York Press.
- Rodriguez-Aguilar, J., Martin, F., Noriega, P., Garcia, P., Sierra, C., 1998. Towards a test-bed for trading agents in electronic auction markets. *AI Communications* 11 (1), 5–19.
- Rodriguez-Aguilar, J., Sierra, C., 2002. Enabling open agent institutions. In: Dautenhahn, K., Bond, A., Canamero, L., Edmonds, B. (Eds.), *Socially Intelligent Agents: Creating relationships with computers and robots*. Kluwer Academic Publishers, pp. 259–266.
- Searle, J., 1969. *Speech Acts*. Cambridge University Press.
- Sergot, M., 2004a. $(C+)^{++}$: An action language for modelling norms and institutions. Tech. Rep. 2004/8, Department of Computing, Imperial College London, retrieved March 6, 2006, from <http://www.doc.ic.ac.uk/research/technicalreports/2004/DTR04-8.pdf>.
- Sergot, M., 2004b. Modelling unreliable and untrustworthy agent behaviour. In: Dunin-Keplicz, B., Jankowski, A., Skowron, A., Szczuka, M. (Eds.), *Proceedings of Workshop on Monitoring, Security, and Rescue Techniques in Multiagent Systems (MSRAS)*. *Advances in Soft Computing*. Springer-Verlag, pp. 161–178.
- Sergot, M., Craven, R., 2006. The deontic component of action language nC+. In: Goble, L., Meyer, J.-J. C. (Eds.), *Deontic Logic and Artificial Normative Systems*. Proc. 8th International Workshop on Deontic Logic in Computer Science (DEON’06), Utrecht, July 2006. LNAI 4048. Springer Verlag, pp. 222–237.
- Shanahan, M., 1999. The event calculus explained. In: Wooldridge, M., Veloso, M. (Eds.), *Artificial Intelligence Today*. LNAI 1600. Springer, pp. 409–430.
- Shanahan, M., 2000. An abductive event calculus planner. *Journal of Logic Programming* 44, 207–239.
- Shoham, Y., Tennenholtz, M., 1992. On the synthesis of useful social laws for artificial agent societies. In: *Proceedings of Conference on Artificial Intelligence*. The AAAI Press/ The MIT Press, pp. 276–281.
- Shoham, Y., Tennenholtz, M., 1995. On social laws for artificial agent societies: off-line design. *Artificial Intelligence* 73 (1-2), 231–252.
- Singh, M., 2000. A social semantics for agent communication languages. In: Dignum, F., Greaves, M. (Eds.), *Issues in Agent Communication*. LNCS 1916. Springer, pp. 31–45.
- Tennenholtz, M., 1995. On computational social laws for dynamic non-homogeneous social structures. *Journal of Experimental and Theoretical Artificial Intelligence* 7.
- Yolum, P., Singh, M., 2002. Flexible protocol specification and execution: applying event calculus planning using commitments. In: Castelfranchi, C., Johnson, L. (Eds.), *Proceedings of Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. ACM Press, pp. 527–535.
- Yolum, P., Singh, M., 2004. 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.