

An Empirical Semantics Approach to Reasoning About Communication

Felix Fischer and Michael Rovatsos

Department of Informatics
Technical University of Munich
85748 Garching, Germany
{rovatsos, fischerf}@cs.tum.edu

Abstract. In the process of designing multiagent systems, it is often the case that some kind of specification of communication rules (in the form of protocols, ACL semantics, etc.) is available. and The question naturally arises how appropriate agents can be designed to can operate on such a specification. Moreover, if these multiagent systems are viewed as *open systems*, the problem is complicated even further by the fact that adherence to such a supposedly agreed specification cannot be ensured on the side of other agents.

This paper presents an architecture for dealing with a very generic type of pre-specified communication patterns (which contain surface structure and logical constraint specifications) based on an *empirical semantics* model of communication. This model allows for flexible adaptation to evolving communication semantics by combining existing expectations about the use of communication with empirical observation.

This architecture is based on the InFFrA social reasoning framework and the concept of interaction frames. We show how interaction frames that represent classes of interaction situations can be used to conduct decision-theoretic reasoning about communication when interpreted using the empirical semantics approach.

We introduce the abstract architecture, a formal model for its probabilistic semantics and present results of an experimental validation of our approach in a complex domain that illustrate its effectiveness.

Keywords Agent interaction, agent communication languages, interaction protocols, multiagent learning, social reasoning

1 Introduction

The process of designing agent communication languages, interaction protocols and conversation policies is primarily concerned with what goes on *between* rather than *inside* agents. It is therefore only natural that ACL research has rarely attacked the problem of how to build agents in accordance with a given set of semantic rules, protocol structures, etc.

In the light of *open* multiagent systems, in which agents are not controlled by a central entity or pursue common goals and need not be benevolent, this design problem

is complicated even further. Not knowing how other agents will apply a given specification concerning the semantic and pragmatic properties of the provided means of communication, the question for an agent is not only how to apply it by himself, but also whether to comply with the supposedly agreed specification at all.

This is exactly the issue this paper deals with: *Given the specification of a number of communication patterns together with logical constraints for their application, how can we design an agent that uses them appropriately to further his own goals, not knowing whether other agents will comply?* As a possible answer to this question, we propose an agent architecture based on the concept of *empirical semantics* [20]. This architecture can be seen as a “communication reasoning” component that takes as inputs the belief state and utility function of an agent as well as a set of models of communication patterns (so-called *interaction frames*) and outputs communication decisions. The architecture is characterised by two essential features:

1. It uses decision-theoretic principles to make optimal action decisions while taking previous communication *experience* into account.
2. It combines empirical knowledge with *a priori* specified logical constraints for communication patterns.

Most importantly, the reasoning heuristics used in this architecture are computationally tractable and can be implemented directly. We thus contribute to the task of engineering agents that are able to operate flexibly on a given set of communication protocols and semantic rules and under the assumption that the empirical semantics view is followed.

The remainder of this paper is structured as follows: In section 2 we review the foundations of the empirical semantics approach. Subsequently, we introduce the InFFrA social reasoning framework that the proposed architecture is based on. Section 4 establishes a formal model of empirical semantics called m^2 InFFrA that uses InFFrA frames to capture probabilistic expectations about communication. In section 5 we then present the apparatus necessary for making decisions in this model. Experimental results from a concrete implementation of the model in a complex domain are given in section 6 to underline its practical usefulness. Section 7 rounds up with some closing remarks and conclusions as well as an outlook to future work.

2 Empirical Semantics

As stated in the introduction, the aim of our research is to develop agents that can reason about communication in *open systems*, in which adherence to a particular communication semantics cannot be taken for granted. That is, no matter what kind of specification is provided for the meaning of utterances used in inter-agent communication, a strict interpretation of agent *autonomy* [12] must include the possibility that agents are entirely self-interested. For such self-interested agents there is no autonomy-respecting way of imposing behavioural constraints, and hence any definition of semantics that includes some *perlocutionary* aspects [1] that affect future actions of an agent (e.g. performing an action one has promised to perform) may be violated in actual communication, if agents decide to deviate from the behavioural pattern supposedly induced by their messages that have been previously exchanged. For this reason, the only way to obtain more

or less reliable models of the relationship between communication and agent action in an entirely open view of communication semantics is through empirical observation and identification of regularities in others' (and one's own) behaviour.

At the same time, it makes little sense to assume that agents act randomly and that the only way to predict their behaviour is to model them as arbitrary stochastic processes. In most systems, some *a priori* agreement about communication semantics exists, whether it be defined in terms of mental states of the participating agents (e.g. [22, 4]) or commitments (e.g. [6, 23]) or in some other way. Also, interaction protocols that are assumed to be common knowledge among agents are usually defined *a priori* so as to restrict the range of admissible message sequences. In other words, from the point of view of a reasoning agent who wants to employ the communication mechanisms provided in a multiagent system, pre-constructed specifications of communication processes and semantics are available, even if their reliability hinges on the ways other agents will make use of the provided specifications.

These considerations lead us to regard pre-specified protocols and ACL semantics as *communicative expectations* rather than *definitions of meaning*, as their properties might be confirmed or violated in actual (future) communication. In the following paragraphs, we explain (i) how a model of meaning can be developed that is in accordance with this expectation-based view and (ii) what the implications of this model are from the standpoint of rational agent-level decision making.

2.1 Expectation-Based view

Our approach is based on the notion of *empirical semantics* for multiagent systems, which was first developed in [20] and [10] (but see also [11, 13, 14]). The empirical semantics view is based on the simple idea that if we assume the meaning of agent communication to be based on expectations (rather than reliable and statically verifiable conditions), we must adapt it according to empirically observed communication and action in the system.

As a consequence, communication must be thought of as a process of exchanged observable symbols which obtain their meaning through the interactions within which they occur. The central instrument to establish this meaning are the expectations constructed by agents participating in or observing an interaction. These expectations “contain” the current semantics of communication in a given social context and their evolution over time mirrors the evolution of meaning.

To make things a bit more concrete, we quote the following central elements of the empirical semantics approach from [20]:

1. The meaning of a message can only be defined in terms of its *consequences*, i.e. the messages and actions that are likely to follow it. Two levels of effects can be distinguished:
 - (a) The immediate reactions of other agents and oneself to the message.
 - (b) The “second-order” impact of the message on the expectation structures of any observer, i.e. the way the utterance alters the causal model of communicative behaviour.




Fig. 1. Empirical semantics and its evolution

2. Any knowledge about the effects of messages must be derived from *empirical* observation. In particular, a semantics of protocols cannot be established without taking into account how the protocols are *used* in practice.
3. Meaning can only be *constructed* through the eyes of an agent involved in the interaction, it strongly relies on relating the ongoing communication to the agent's own goals.

In other words, our view is that the semantics of communication can be reduced to its pragmatics as perceived by agents who are using communication. While this is an overtly restrictive interpretation of meaning for general communication, it can be very useful for practical agent reasoning in *open* multiagent systems, in which “normative” semantics may be proven wrong by others’ inability or unwillingness to comply with them.

This is because the empirical semantics view can be used as a *minimal-assumption* method for reasoning about communication which at least ensures that the agent will not compromise his own welfare for the sake of complying with a pre-specified semantics (when others do not comply, either). In other words, we can always “resort to” observed regularities between utterances and subsequent (environment-manipulating, i.e. “physical”) actions if all else fails.

Taking this as a point of departure for reasoning about communication raises the following question: If the meaning of an utterance (or sequence of utterances) lies in the expected consequences, how can we capture semantics using a model of causality and correlation that allows to predict these consequences for a given (sequence of) utterance(s)? What is foremostly required to solve this problem are appropriate representations for adaptable communicative expectations, i.e. representations that reflect expectations and that can be adapted with new observations so as to enable the agent to extrapolate past correlations into the future.

Expectation networks (ENs, [9, 13]) have been suggested as a very general method of capturing this kind of expectations that relies on a probabilistic model of “continuations” between interrelated messages derived from statistical observation. Instead of going into the details of this formalism, we sketch the basic idea using the illustrations of figure 1. In this figure, nodes represent communicative actions, edges correlations between them (variable line width is used to indicate different degrees of correlation). The shaded node sequence is used to describe the recently observed portion of the net-

work, e.g. an ongoing conversation. The decreasingly dark shaded regions of predicted future actions denote that predictions regarding “distant” events are increasingly vague. The transition from the situation shown on the left to that on the right occurs upon observation of a new action that is appended to the currently relevant path. With this new observation, the correlation between the message previously observed and the current message increases compared to alternatives that did not occur.

This simple probabilistic process modelling view is obviously inadequately simplistic and suggests that agents are nothing but the aforementioned stochastic processes whose behaviour is arbitrary and not confined in any way by principles of rationality, goal orientation or logical inference. In fact, this is quite deliberate for cases in which other agents do not seem to be behaving according to any “higher-level” models the observer may have, i.e. if the assumptions he makes regarding other agents, their mental states, the state of commitments in the system etc. and – in particular – that part of these assumptions that is concerned with communication semantics and interaction rules fails. Hence, the purely probabilistic view is only useful in a zero-assumption case in which all an observer can do is to correlate message and action sequences.

Yet the expressiveness of ENs goes far beyond that, as they contain elements that we have not described in our oversimplified description of the EN formalism above: Firstly, instead of messages agents actually use message patterns that may involve variables as node labels, so that a path can be a generalisation of a set of observed ground message sequences. Additionally, sets of instance values for these variables can be stored in the node to re-construct previous cases generalised over if necessary (or to enhance prediction capabilities with the help of similarity considerations for new cases in a case-based reasoning fashion). Secondly, transitions between nodes are not only labelled with numerical probability estimates derived from past frequency data, but may also contain logical constraints that restrict the applicability of certain transitions.

With these additional facilities, ENs become a highly powerful tool to add semantic information to expectations that reflects assumptions the observer makes, thus allowing for a flexible mix between hard-coded constraints and adaptive probabilistic reasoning.

2.2 Decision-Theoretic view

In principle, adopting the standpoint of empirical semantics enables an agent to adapt his expectations with new communication experience and to adjust his own behaviour to the expected reactions of others. However, this conceptual view of semantics does not say anything about how agents can use such expectations in practice.

What is therefore required is to integrate expectations in the decision-making processes of an agent, and to this end we shall discuss a few decision-theoretic considerations that are related to communication in different interaction situations.

As two famous examples for strategic interaction situations, consider the Prisoner’s Dilemma (PD) game and the Coordination Game (CG), which are normal-form two-player games characterised by the payoff matrices shown in tables 1 and 2.

In the PD game, there is an incentive for both agents to defect, since D is a better choice whatever the other party chooses (hence (D, D) is the only Nash equilibrium in this game). The dilemma arises because agents could both do better under the (Pareto

	a_j	C	D
a_i			
C		(3,3)	(0,5)
D		(5,0)	(1,1)

Table 1. Prisoner’s Dilemma payoff matrix. Matrix entries (u_i, u_j) contain the payoff values for agents a_i and a_j for a given combination of row/column action choices, respectively. C stands for each player’s “cooperate” option, D stands for “defect”.

	a_j	A	B
a_i			
A		(1,1)	(-1,-1)
B		(-1,-1)	(1,1)

Table 2. Coordination game, in which agents have two choices A and B and receive a positive payoff if they opt for the same alternative.

efficient) (C, C) combination, which neither of them is probably going to choose because it may imply getting the “sucker’s payoff” 0 if the other defects. So the problem is that while agents’ interests overlap in this game, cooperation is impeded by the potential for mutual exploitation. In the CG, on the other hand, reaching an agreement is trivial in theory, since agents receive the equal rewards. However, in the absence of any means of communication, no agent can know what the other will choose, resulting in a (mixed-strategy Nash equilibrium) solution for the game where each agent will play A or B with equal probability and only obtain an expected payoff of 0.

Now let us assume prior to actually playing any of these games agents may exchange a number of messages. According to our consequentialist view, the only way to interpret these messages would be as indicators for the “physical” actions $(C/D$ or $A/B)$ that they will perform in the actual game. This is also the only realistic significance messages would have in such an interaction scenario from a game-theoretic point of view if the interaction situation is completely determined by the definitions of the one-shot games (i.e. no externalities have to be taken into account to which communication might contribute).

In the PD, obviously, it makes no sense for any one agent to indicate his choice, because he might be exploited if he is telling the truth, and the other agent has nothing to lose by defecting (he can only gain from exploiting the other, should he be truthful). In the CG game, on the other hand, truthfully indicating what one will do is a dominant strategy for both agents, and it makes no sense for them to lie.

This illustrates the different flavours communication may take on, ranging from self-enforcing communication sharing (as in the CG) that merely serves to synchronise agent activity to highly contingent exchange of proposals, to (potentially fraudulent) offers for cooperation that are highly contingent and depend on the other’s cooperative attitude.

But what is the significance of such communication from an empirical semantics point of view? Under a consequentialist view of meaning, what we can infer from the above examples is that any symbol uttered by an agent represents a course of joint future actions (or, a set of different alternatives) in an interaction situation. Given an EN-like representation of expectations, any message that corresponds to a node in the network and is being uttered in the context of an ancestor path (e.g. by restricting the temporal scope of what counts as “relevant” to the current interaction situation) denotes the sum of its descendant paths (again, potentially restricted in depth and breadth by some notion of interaction scope). So depending on the structure of the EN (which may also suggest that certain paths are valid or invalid at the time of utterance depending on logical constraints or may induce a stochastic distribution over possible continuations in a probabilistic fashion) the agent uttering a message is selecting a set of possible continuations, thereby “requesting” a (set of) desired outcome(s) from the other(s). The message obtains precisely the predicted set of possible continuations of the current message (viz outcomes of the interaction situation) as its intended meaning. Under this interpretation, any message is a description of a set of expected outcomes as anticipated by its utterer, and if we trace a sequence of messages, we can view subsequent steps in this sequence as an incremental refinement of requests (or offers) for possible states of the world after the interaction. Along this communicative “encounter”, the agents make communication decisions according to their subjective expectations with the goal of reaching states that are desirable for themselves. If expectations are mutually compatible in the sense that the agents understand what each other is saying and trusts that the expected outcomes will actually occur, and if there is sufficient overlap between agents’ interests, reliable expectation structures can lead to coordinated behaviour.

As an illustration, let us look at an example of an expectation structure in which an agent may indicate he wants to play a PD game by uttering an initial message, wherupon the other party may agree or disagree to play. After this initial stage, the agents may choose from different alternative messages to indicate what they will do, and the eventually occurring actual game moves are performed concurrently. This situation is shown in figure 2. Depending on the probabilities of all transitions (and especially of those

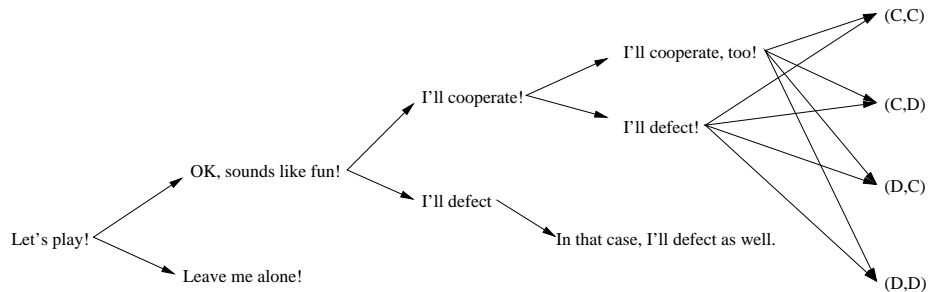


Fig. 2. Expectation structure for pre-play communication in the PD game; joint actions occur concurrently, while communication steps alternate between the two agents.

leading to leaf nodes) and the utilities assigned to leaf nodes which are omitted in the figure, agents can calculate best-/worst-case or expected utilities depending on their decision rules so as to make communication and action choices at each step.

At a more practical level, if communication is inexpensive in comparison to the outcomes of the physical actions, this leads to a “cheap talk” [2] view of communication, in which agents engage in conversations to maximise their own profit under the assumptions of their own EN. Obviously, the usefulness of this EN depends strongly on its reliability, i.e. on how well the symbol meaning is common knowledge and adhered to. However, things can go wrong if agents do not share expectations or are deliberately exhibiting *deviant* behaviour (we refer the interested reader to [18] for a detailed discussion of deviance and, as one of its most interesting manifestations i.e. conflict).

To bridge the gap between the expectation-based and the decision-theoretic views, we need to devise appropriate representations that agents can use to record and reason about their interaction experience. For interactions in the form of two-party, turn-taking dialogues, we suggest *interaction frames* as one such representation. The following sections introduce the InFFrA architecture that is used for this purpose and discuss the reasoning mechanisms employed in .

3 The InFFrA Architecture

The Interaction Frames and Framing Architecture InFFrA [21] has originally been developed as a meta-framework for social reasoning architectures based on the notions of “interaction frames” and “framing”. It provides abstractions for the data and processes involved in reasoning about patterns of communication at a meta-architectural level, i.e. without pinning down concrete implementation details. In the following, a brief introduction to the abstract framework is provided. After this, we will explain how this framework can be combined with the empirical semantics view laid out in the previous sections.

3.1 Abstract Architecture

The central idea behind InFFrA is to employ models of classes of interaction called interaction frames to guide agents’ social behaviour. The process of applying frames appropriately in interaction situations is referred to as framing. In the abstract InFFrA architecture, a frame is a data structure that contains information about¹

- the possible courses of interaction (so-called *trajectories*) characteristic to a particular frame,
- *roles and relationships* between the parties involved in an interaction of this class,
- *contexts* within which the interaction may take place, and
- *beliefs*, i.e. epistemic states of the interacting parties.

¹ We include only so-called *descriptive* frame attributes in our discussion. In the actual model [18], frames may also contain *meta-level* attributes that describe relationships between frames, frame history, statements about the extent to which knowledge about frames is distributed among agents, etc.

At the conceptual level, this allows for modelling all relevant features of a class of interactions: its participants, its surface structure, the context within which instances of it may occur and the associated beliefs of participating parties. InFFrA makes no assumptions as regards representation, thus also allowing for informal and semi-formal modelling methods where appropriate.

For computational purposes, a representation of a set of admissible message and action sequences is usually taken to represent the trajectory model (e.g. a DFA, a Petri Net, etc.) while the latter three elements can be collapsed into a single set of logical constraints whose fulfillment has to be verified using the agent's internal belief state (usually represented by some kind of knowledge base for which tractable proof procedures are available) if the respective frame is to be considered applicable in a given situation.

InFFrA makes use of a number of frame-based data structures to conduct the steps necessary for framing:

- The *active frame*, the unique frame currently activated to describe the expected course of events,
- the *perceived frame*, an interpretation of the currently observed state of affairs,
- the *difference model* containing the differences between perceived frame and active frame,
- the *trial frame*, used when alternatives to the current frame are sought for,
- and the *frame repository*, in which the agent locally stores its frame knowledge.

Using these data structures, an InFFrA agent performs the following steps in each reasoning cycle:

1. *Interpretation & Matching*: Update the perceived frame and compare it with the active frame.
2. *Assessment*: Assess the usability of the active frame in terms of
 - (i) adequacy (compliance of frame conditions with the current situation),
 - (ii) validity (the degree to which the active frame's trajectory matches the perceived encounter) and
 - (iii) desirability (depending on whether the implications of the frame correspond to the agent's private goals).
3. *Framing decision*: If the active frame seems appropriate, continue with 5. Else, proceed with 4 to find suitable alternatives.
4. *Adjustment/Re-framing*: Search the frame repository for better frames. "Mock-activate" them as trial frames iteratively and go back to 1; if no suitable frame is found, end the encounter.
5. *Enactment*: Derive action decisions by applying the active frame.

The entire² framing process is depicted in figure 3. It should be emphasised that InFFrA only describes the social layer of agent reasoning; in order to obtain an integrated

² This is a largely simplified view of the actual InFFrA-based social reasoning cycle. [18] introduces a much more intricate data flow model that involves complex interactions between the data structures, a cyclic model of trial instantiation, distinguishing between frame enactment and actual behaviour generation, etc.

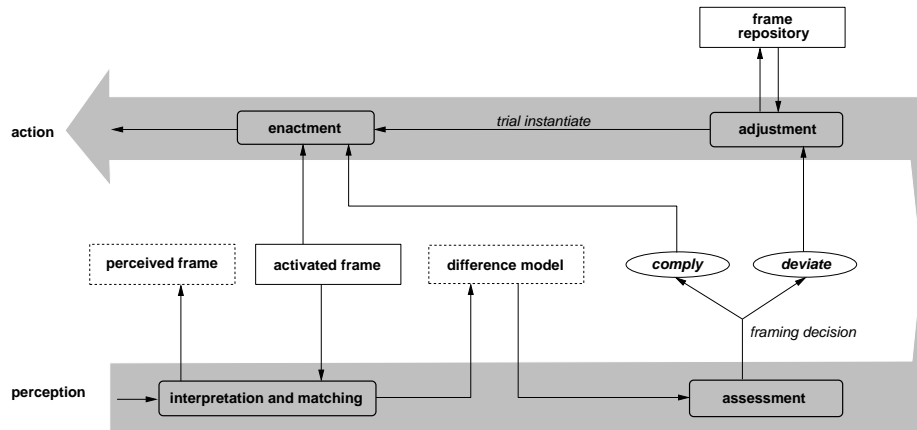


Fig. 3. Overview of the framing process

agent architecture, it has to be combined with a suitable component for local rational reasoning (e.g. a BDI [16] reasoner), so that the agent's private goals and preferences can be taken into consideration during the *assessment* phase.

3.2 Frames as Expectation Structures

Looking at InFFrA from an empirical semantics point of view, we can re-interpret the frames in an agent's repository as the expectations about communication patterns that the agent holds.

In *retrospective* terms, the agent can use observed interaction situations to derive new frames (or store them as instances of already known frames) and extend his own repository using new interaction experience. This will allow to re-construct interaction experience (ideally) from a reasonably small set of classes of interactions. In *prospective* terms, the agent can apply his repository to determine a set of admissible frames in any given interaction situation and to use these for prediction of possible continuations. This enables the application of decision-theoretic principles to frame selection and frame adjustment (and, in very elaborate architectures also to creative construction of completely novel frames).

The frame-based approach lends itself to reasoning with empirical semantics for two main reasons:

- InFFrA provides us with a practical reasoning scheme for processing expectations which is missing in the pure expectation-based view of expectation networks described above.
- InFFrA is ideally suited for combining pre-specified communication patterns (e.g. protocols and ACL rules specified by the designer) in a goal-oriented fashion with empirical knowledge about how the interaction actually *turns out* in the system.

In the next section, we are going to explain how a formal model of an instance of the InFFrA architecture can be developed that is in keeping with these considerations.

4 A Formal Model of Frame-Based Empirical Semantics

To obtain a formal computational model for representing and reasoning about communicative expectations based on InFFrA, we have developed m^2 InFFrA, a simple yet expressive model for describing two-party, discrete, turn-taking *encounters* which can be thought of as conversations between two agents.

4.1 The m^2 inffra model

m^2 InFFrA frames use with linear sequence of message patterns (i.e. messages containing variables) as its trajectory model specifying the surface structure of the encounters described by a particular frame, together with a list of *substitutions* to capture concrete values of these variables in previously experienced interactions. Each substitution also corresponds to a set of logical *conditions* that were required for and/or precipitated by execution of the trajectory in the respective encounter. Finally, *trajectory occurrence* and *substitution occurrence* counters record the frequency with which the frame has occurred in the past. Formally, a frame in m^2 InFFrA is defined as follows:

Definition 1. A frame is a tuple $F = (T, \Theta, C, h_T, h_\Theta)$, where

- $T = \langle p_1, p_2, \dots, p_n \rangle$ is a sequence of message patterns $p_i \in \mathcal{M}$, the trajectory of the frame,
- $\Theta = \langle \vartheta_1, \dots, \vartheta_m \rangle$ is an ordered list of variable substitutions,
- $C = \langle c_1, \dots, c_m \rangle$ is an ordered list of condition sets, such that $c_j \in 2^{\mathcal{L}}$ is the condition set relevant under substitution ϑ_j ,
- $h_T \in \mathbb{N}^{|T|}$ is a trajectory occurrence counter list counting the occurrence of each prefix of the trajectory T in previous encounters, and
- $h_\Theta \in \mathbb{N}^{|\Theta|}$ is a substitution occurrence counter list counting the occurrence of each member of the substitution list Θ in previous encounters.

In this definition, \mathcal{M} is a language of speech-act [1] like message and action patterns of the form $\text{perf}(A, B, X)$ or $\text{do}(A, Ac)$. In the case of messages (i.e. exchanged textual signals), perf is a performative symbol (`request`, `inform`, etc.), A and B are agent identifiers or agent variables and X is the propositional content of the message taken from a logical language \mathcal{L} . In the case of physical actions (i.e. actions that manipulate the physical environment) with the special “performative” `do`, Ac is the action executed by A (a physical action has no recipient as it is assumed to be observable by any agent in the system). Both X and Ac may contain non-logical substitution variables that are used for generalisation purposes (as opposed to logical “content variables” used by agents to indicate quantification or to ask for a valid binding). We further use $\mathcal{M}_c \subset \mathcal{M}$ to denote the language of “concrete” messages that agents use in communication (and that do not contain variables other than “content variables”).

To illustrate these concepts and to explain the semantics of a frame, we will consider an example of how the FIPA contract net protocol [5] can be implemented in m^2 InFFrA.

$$\begin{aligned}
F_{cn} = & \left\langle \left\langle \xrightarrow{5} \text{cfp}(A_1, A_2, \langle R, P \rangle) \xrightarrow{3} \text{propose}(A_2, A_1, Q) \right. \right. \\
& \xrightarrow{3} \text{accept-proposal}(A_1, A_2, Q) \xrightarrow{2} \text{do}(A_2, A_1, R) \left. \right\rangle, \\
& \left\langle \{ \iota X (P = Q) = Y, \right. \\
& \quad \neg B_{ref_{A_1}}(\text{any } X I_{A_2} \text{Done}(R, P)) \wedge \neg B_{A_1} I_{A_2} \text{Done}(R) @1, \\
& \quad B_{A_2} I_{A_2} \text{Done}(R, Q) @2, \\
& \quad B_{A_1} I_{A_1} \text{Done}(R, Q) \wedge B_{A_1} I_{A_2} \text{Done}(R, Q) @3, \\
& \quad B_{A_2} Q @4 \left. \right\rangle, \{ \}, \{ \text{damaged}(\text{tires}) \} \left. \right\rangle, \\
& \left\langle \xrightarrow{0} [] \right. \\
& \quad \xrightarrow{1} [A_1/\text{agent}_1, A_2/\text{agent}_2, P/\text{priceOf}(\text{tire}, X), \\
& \quad \quad Q/\text{priceOf}(\text{tire}, 75), R/\text{sell}(\text{tire}, 4)], \\
& \quad \xrightarrow{1} [A_1/\text{agent}_3, A_2/\text{agent}_1, P/\text{priceOf}(\text{tire}, X), \\
& \quad \quad Q/\text{priceOf}(\text{tires}, 400), R/\text{replace}(\text{tires})] \left. \right\rangle,
\end{aligned}$$

Table 3. Interaction frame for the success path of the FIPA contract net protocol

For ease of presentation, we will write $T(F)$, $C(F)$, etc. to denote the respective elements of a frame F and use the compact notation $\langle T_h(F), C(F), \Theta_h(F) \rangle$ instead of $(T, C, \Theta, h_T, h_\Theta)$, where

$$T_h(F) = \xrightarrow{h_T(F)[1]} p_1 \xrightarrow{h_T(F)[2]} p_2 \cdots \xrightarrow{h_T(F)[n]} p_n$$

and

$$\Theta_h(F)[i] = \xrightarrow{h_\Theta(F)[i]} \Theta(F)[i].$$

Table 3 shows an interaction frame for the success path of the contract net protocol, the following should be noted about this example:

- As can be seen, the first condition set (corresponding to the empty substitution) contains feasibility preconditions of the respective performatives (in FIPA-SL [4] with additional timestamps of the form @ i indicating at which trajectory step a condition has to hold), as far as they are relevant for frame execution.
- Definition 1 does not constrain which conditions are to be stored for a specific enactment of a frame (even allowing for empty condition sets), and the task of extracting *relevant* or even *crucial* information from the agent’s knowledge is clearly nontrivial. However, the reasoning framework to be defined in the following section primarily uses conditions to identify similarities in encounters, while expectation is drawn from utilities obtained during frame execution, making this approach less sensitive to the way conditions are selected.

As for the different failure cases covered by the contract net protocol, these could either be modelled implicitly by using timeouts and interpreting “silence as disapproval” or explicitly by virtue of additional frames. For example, the frame of table 4 models the refusal to submit a proposal.

$$\begin{aligned}
F_{cr} = & \langle \langle \xrightarrow{5} \text{cfp}(A_1, A_2, \langle R, P \rangle) \xrightarrow{2} \text{refuse}(A_2, A_1, Q) \rangle, \\
& \langle \{ \neg B_{ref} A_1 (\text{any } X I_{A_2} \text{Done}(R, P)) \wedge \neg B_{A_1} I_{A_2} \text{Done}(R) @1, \\
& B_{A_2} (\neg \text{Feasible}(R) \wedge Q \wedge \neg \text{Done}(R) \wedge \neg I_{A_2} \text{Done}(R)) @2 \}, \\
& \{ B_{A_1} (\text{inStock}(\text{tire})) @1 \} \rangle, \\
& \langle \xrightarrow{0} [] \rangle \\
& \xrightarrow{2} [A_1/\text{agent}_1, P/\text{priceOf}(\text{tire}, X), \\
& Q/\neg \text{inStock}(\text{tire}), R/\text{sell}(\text{tire}, 4)]
\end{aligned}$$

Table 4. Interaction frame modelling the refusal to submit a proposal in the FIPA contract net protocol

Together, F_{cn} and F_{cr} capture the following observations about previous encounters:

- Five encounters started with a message matching $\text{cfp}(A_1, A_2, \langle R, P \rangle)$, three of them continued with a proposal by A_2 , the other two with a refusal. In two of the former three cases, the proposal was accepted by A_1 and the respective physical action carried out by A_2 .
- One encounter has terminated after the second message or was continued with a message not matching $\text{accept-proposal}(A_1, A_2, Q)$.
- For four of the five encounters, substitutions are available. Two contracts were made, one about the delivery of four tires for a price of 75 each, one about the replacement of a set of tires for 400. Two calls for proposals for the delivery of four tires were refused because no tires were available (which is a generalisation over two encounters, leaving A_2 unspecified).

4.2 Making predictions with m²inffra frames

To use m²InFFrA frames for the prediction of future encounters we need a formal model of their *prospective* rather than the *retrospective* semantics informally described in the previous section.

Here, the idea is (as in [20]) to use an entire repository $\mathcal{F} = \{F_1, \dots, F_n\}$ of frames, each representing a set of message/action sequences by virtue of the substitutions that can still be applied to its trajectory. Given an *encounter prefix*, i.e. a sequence of messages already uttered in the current encounter, and the agent's current belief state, we can filter out those paths that either (a) do not match the encounter prefix or (b) are labelled with logical conditions not satisfied under current knowledge base content. Considering the remaining (i.e. relevant) paths, we can then assign probabilities to all the possible *encounter continuations* (or postfixes) using the counter values of $F \in \mathcal{F}$. Additionally, domain-dependent *case-based reasoning* [8] techniques can be applied by introducing a similarity measure on messages, such that different instances of a postfix pattern have different probabilities depending on their similarity with past enactments

of the respective frame. Hence, \mathcal{F} represents a simplified version of an expectation network that has the form of a tree.

More formally, let $w \in \mathcal{M}_c^*$ the encounter prefix and

$$\vartheta_{fixed}(F, w) = \text{unifier}(w, T(F)[1:|w|])$$

the most general unifier (MGU) of w and the corresponding trajectory prefix $T(F)[1:|w|]$ of F (where $\text{unifier}(\cdot, \cdot)$ returns the most general unifier for two message patterns or sequences thereof, or \perp if they cannot be unified).

For a given knowledge base $KB \in 2^{\mathcal{L}}$ describing the belief state of an agent (KB is assumed to be encoded in the same propositional language \mathcal{L} as the content of messages for reasons of simplicity), this allows us to define

$$\Theta_{poss}(F, KB, w) = \{\vartheta \mid \exists \vartheta'. \vartheta = \vartheta_{fixed}(F, w)\vartheta' \wedge \exists i. KB \models C[i]\vartheta\}$$

as the set of substitutions still *possible* under F , KB , and w . The elements of Θ_{poss} are extensions of ϑ_{fixed} for which at least one condition in $C(F)$ is satisfied.³

For a given *similarity measure* $\sigma : \mathcal{M}^* \times \mathcal{M}^* \rightarrow \mathbb{R}$ on message pattern sequences, we can define

$$\sigma(\vartheta, F) = \sum_{i=1}^{|\Theta(F)|} \overbrace{\sigma(T(F)\vartheta, T(F)\Theta(F)[i])}^{\text{similarity}} \overbrace{h_{\Theta}(F)[i]}^{\text{frequency}} \overbrace{c_i(F, \vartheta, KB)}^{\text{relevance}}$$

to assess to which extent ϑ is “applicable” to F . In this definition, $h_{\Theta}(F)[i]$ is used to take the frequency of a past condition/substitution into account and c_i expresses how relevant a particular frame condition $C[i]$ is in determining the applicability of ϑ .⁴ This quantity can be used to derive a conditional probability distribution over the different substitutions ϑ F may be enacted under:

$$P(\vartheta|F, w) = \begin{cases} \alpha \cdot \sigma(\vartheta, F) & \text{if } \vartheta \in \Theta_{poss}(F, KB, w) \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

for some normalisation constant α . Finally, combining this formula with the frequency

$$P(F|w) = \begin{cases} \frac{h_T(F)[|T(F)|]}{\sum_{F' \in \mathcal{F}, \text{unify}(T(F')[1:|w|], w)} h_T(F')[|T(F')|]} & \text{if } \text{unify}(T(F)[1:|w|], w) \\ 0 & \text{else} \end{cases}$$

³ We use the notation $L\vartheta$ here for the result of applying ϑ to each element of a list or set L , and $\vartheta\vartheta'$ for the substitution that results from applying ϑ' after ϑ . Further, we implicitly assume that only *minimal* substitutions are considered for ϑ' , only replacing variables that actually occur in $T(F)$ or in some $C(F)[i]$.

⁴ A simple definition of c_i , considering only those conditions currently satisfied, is for example given by

$$c_i(F, \vartheta, KB) = \begin{cases} 1 & \text{if } KB \models C(F)[i]\Theta(F)[i]\vartheta \\ 0 & \text{otherwise.} \end{cases}$$

of F matching any past encounter starting with w allows us to compute the *continuation probability* with which an encounter that started with w will be concluded with w' :

$$P(w'|w) = \sum_{F \in \mathcal{F}, ww' = T(F)\vartheta} P(\vartheta|F, w)P(F|w)$$

Looking back at figure 1, this equation defines a probability distribution for the possible continuations of an encounter given a current “path” in the expectation network induced by \mathcal{F} . Next, we will show how this semantics can be used to conduct rational reasoning about communication in $m^2\text{InFFrA}$.

5 Reasoning About Communication in $m^2\text{inffra}$

Based on the formal model presented in the previous section, the general principles of InFFrA can be concretised and embedded into an agent architecture to endow agents with the ability to reason and make decisions about communication. To illustrate how this is done, we will again go through the individual steps of the abstract reasoning cycle depicted in figure 3.

Interpretation and Matching/Assessment At the beginning of each reasoning cycle, the knowledge base KB and the encounter prefix w are updated from the peer’s last utterance.

As for matching and assessment, the agent checks if $\Theta_{\text{poss}}(F_a, KB, w) \neq \emptyset$, i.e. if the trajectory $T(F_a)$ of the active frame F_a matches w and the remaining steps of $T(F_a)$ can still be executed under KB . Considerations about the *desirability* of F_a could also play a role in the framing decision, and a possible definition for such a desirability criterion will be given in our description of frame enactment.

Returning to the example of the previous section, we consider an encounter with $\mathcal{F} = \{F_{cn}, F_{cr}\}$ and $F_a = F_{cn}$ for both agents and prefix $w = \text{cfp}(\text{agent}_2, \text{agent}_1, (\text{sell}(\text{tire}, 1), \text{priceOf}(\text{tire}, X)))$, so that agent_1 can successfully match F_a against w . If for example agent_1 did only sell tires in fours, executability of F_a would still fail and require a re-framing.

Adjustment/Re-Framing The idea behind grouping different courses of interaction into frames is to exploit the fact that (usually) similar types of interaction (e.g. negotiation dialogues, contracting, etc.) exist which differ only in the specific content of messages, but not in what they achieve for the interacting parties. This hierarchical view enables agents to optimise *within* the current frame while disregarding other frames during a conversation. However, if the active frame cannot be carried out any longer, the search space for appropriate continuations has to be expanded.

In this case, a variety of frame selection heuristics can be used to find suitable alternatives. In our implementation, experimental results for which are presented in the following section, we use hierarchical reinforcement learning techniques [24] to learn an optimal frame selection strategy over time. Also, we construct new frames through concatenation in a planning-like manner to achieve the original goal of a conversation that went awry. An extensive treatment of the resulting architecture can be found in [3].

At the end of an encounter – in order to maintain a concise model of past interactions – the active frame is augmented by a substitution that unifies its trajectory with w , along with a set of conditions that were required for or precipitated by the execution.

Enactment If the active frame contains no further steps, the agent simply terminates the encounter, as is the case if no active frame could be found. Else, the next message or action is chosen by applying the locally optimal substitution ϑ^* to the next step of the active frame’s trajectory.

To determine ϑ^* , we assume that the agent’s preference towards different world states is expressed by means of a real-valued *utility function* $u : \mathcal{M}^* \times 2^{\mathcal{L}} \rightarrow \mathbb{R}$, where $u(w, KB)$ is the utility associated with a message sequence w being executed for initial knowledge base KB .

Since some variables of ϑ^* will be bound by the agent himself, while concrete values for others will be “selected” by his peer with a certain conditional probability, the optimal substitution is defined as the one with the highest *expected utility*. Normally, this notion of expected utility will also be used during frame assessment to determine the desirability of a specific substitution or entire frame. For example, one could force a re-framing unless executing the postfix of the active frame under the most desirable substitution yields a positive utility.

If we write Θ_s and Θ_p for the sets of possible substitutions the agent and his peer in the current encounter can apply, respectively, the expected utility of executing a frame F under $\vartheta_s \in \Theta_s$ is given by

$$E[u(\vartheta_s|F, w, KB)] = \sum_{\vartheta_p \in \Theta_p} u(\text{postfix}(T(F), w)\vartheta_s\vartheta_p, KB) \cdot P(\vartheta_p|\vartheta_s, F, w),$$

where $\text{postfix}(T(F), w)$ is the postfix of F corresponding to prefix w (which can be determined by applying the most general unifier of w and the corresponding prefix of $T(F)$ to the respective rest) and $P(\vartheta_p|\vartheta_s, F, w)$ is the probability with which the peer will conditionally choose some substitution $\vartheta_p \in \Theta_p$ depending on the agent’s choice $\vartheta_s \in \Theta_s$.

An approximation of $P(\vartheta_p|\vartheta_s, F, w)$ can be computed from the past cases stored in F . Using Bayes’ rule and applying equation 1 to both numerator and denominator, we can approximate

$$P(\vartheta_p|\vartheta_s, F, w) = \frac{\sigma(\vartheta_{\text{fixed}}(F, w)\vartheta_s\vartheta_p, F)}{\sum_{\vartheta} \sigma(\vartheta_{\text{fixed}}(F, w)\vartheta_s\vartheta, F)}$$

if $\vartheta_{\text{fixed}}(F, w)\vartheta_s\vartheta_p \in \Theta_{\text{poss}}(F, KB, w)$ (and 0, else). By means of standard expected utility maximisation, the optimal substitution ϑ^* is then given by

$$\vartheta^*(F, w, KB) = \arg \max_{\vartheta_s \in \Theta_s} E[u(\vartheta_s|F, w, KB)].$$

Concluding the above example under the assumption that no re-framing was necessary, *agent*₁ now has to select an appropriate proposal. That is, each of the elements of Θ_{poss} yields a different price for a tire, and the search for the best substitution is a search for the highest possible price (as u will usually increase with higher profit) such

that $agent_2$ will still accept (as the probability for the accept will usually decrease with an increasing price) and will be guided by similar past cases stored in $\Theta(F_{cn})$, i.e. by past prices for tires.

6 Experimental validation

In order to show their performance in practice, the concepts presented so far have been implemented and tested in the multiagent-based link exchange simulation system **LIESON** [17]. In this system, agents representing Web sites engage in communication to negotiate over mutual linkage with the end of increasing the popularity of one's own site and that of other preferred sites. Prior to reporting on the actual simulation results obtained with simple proposal-based negotiation frames in this application scenario, we briefly discuss the **LIESON** system and how $m^2InFFrA$ agents are realised in this simulation system.

6.1 The **LIESON** system

The **LIESON** environment is populated by a number of agent who represent web sites and seek to optimise those web-sites linkage situation on behalf of human web site owners. It is assumed that in-links increase traffic on a site, and that agents are trying to maximise the popularity of their own site, and that of other sites which they rate highly. Private ratings are not visible to other agents and represent the site owner's opinion of the content of other sites. In contrast to private ratings, numerical link weights can be used to display *public* ratings that can be observed by anyone in the system.

The only available physical (i.e. utility-affecting) actions in this domain are the addition and deletion of numerically rated links originating from one's own site and the modification of rating values (where the probability of attracting more traffic through a link depends on the rating value). Additionally, agents may engage in low-cost communication with each other to request certain physical actions of others.

Thus, the strategic aspect of this application derives from the fact that agents can use public weights to influence others' popularity and that they can employ this "power" to influence their peer's behaviours. At the same time, there is a dilemma between being honest about one's own opinion of others (after all, one aims to "speak his mind") and agreeing to deviate from private belief in selecting public ratings in order to "massage" the other into a more cooperative stance.

LIESON provides a highly dynamic and complex interaction testbed for the following reasons:

- Agents only have a partial and incomplete view of the link network. In particular, agents engage in non-communicative goal-oriented action in between encounters, so that the link network (and hence the agents' utility situation) may change while a conversation is unfolding.
- The number of possible link configurations is vast, and agents can only predict possible utilities for a very limited number of hypothetical future network layouts.

- There is no notion of commitment – agents choose frames in a self-interested way and may or may not execute the physical actions that result from them. Also, they may undo the effects of these actions later on.

As utility benchmarks for our experiments we use the average score⁵ of all agents in two particular link network layouts:

- *Complete, honest linkage*: The network is fully connected and each link displays the truthful private rating the source site agent holds of the target site.
- *Complete, politically correct rating*: The network contains exactly those links whose private rating values are positive, i.e. if an agent has a positive opinion of another agents the two sites will be connected by a link displaying the correct private rating, if not the source site agent “remains silent” regarding hisd opinion of the target site.

Interestingly, the politically correct linkage patterns results in a higher average score (and higher social welfare) and, although we have not derived analytical bounds for this utility function, no other linkage configuration seems to be more rational at a global level. In particular, extreme configurations (such as a network with empty link set or fully connected networks with random/highly positive/highly negative link weights regardless of private ratings) yield much smaller utility values. Hence, attaining the level of these benchmark values would be an impressive achievement for agents given that they know nothing about this behaviour of the global utility function.

6.2 m²inffra agents in LIESON

LIESON agents consist of a non-social BDI [16] reasoning kernel that projects future link network configurations and prioritises goals according to utility considerations. If these goals involve actions that have to be executed by other agents (and the agent is not already engaged in another conversation), the m²InFFrA component initiates a new framing process.

This framing process begins with an initial message, by which the agent contacts any agent that can perofm the desired action (in the case of LIESON, the appropriate communication partner can be unambiguously determined since only an the agent owning a web-site can modify the outgoing links of that site) and runs until either (i) the goal in question has been achieved and the encounter can be terminated, (ii) no adequate frame can be found and the agent chooses to terminate the conversation, or (iii) the other party terminates the encounter.

⁵ The score of agents in each rounds is calculated by combining the current popularity of their own site and the popularities of preferred sites, while also taking into account to which extent these sites express similar opinions (in terms of public rating values i.e. link weights). Popularity estimates for hypothetical link configurations are obtained by using a shortest-path model of stochastically behaving web users, where link transition probabilities depend on the numerical link weights.

We report on experiments in which agents were equipped with simple proposal-based frames with the following six trajectory models:

```

request(A, B, X) → accept(B, A, X) → confirm(A, B, X) → do(B, X)
request(A, B, X) → propose(B, A, Y) → accept(A, B, Y) → do(B, Y)
request(A, B, X) → propose-also(B, A, Y) → accept(A, B, Y) →
    do(B, X) → do(A, Y)
request(A, B, X) → reject(B, A, X)
request(A, B, X) → propose(B, A, Y) → reject(B, A, Y)
request(A, B, X) → propose-also(B, A, Y) → reject(B, A, Y)

```

The first three frames allow for accepting to perform a requested action X , making a counter-proposal in which Y is suggested instead of X , or using `propose-also` to suggest that B will execute X if A agrees to execute Y . The last three frames can be used to explicitly `reject` a request or proposal. In that, X and Y are link modification actions; each message is available in every state and incurs a cost that is almost negligible compared to the utilities gained or lost through linkage actions (yet high enough to ensure no conversation goes on forever). Also, agents can always send a `stop` action to indicate that they terminate an encounter if they cannot find a suitable frame.

After their termination, encounters are stored in the frame from which they have originated. For example, agent a_1 would store the encounter `request($a_1, a_2, add(a_2, a_1, 2)$) → reject($a_1, a_2, add(a_2, a_1, 2)$)` by adding a substitution $[A/a_1, B/a_2, X/add(a_2, a_1, 2)]$ to the respective frame together with an automatically generated list of conditions that were required for physical action execution.

As state abstraction, we use generalised lists of statements of the form $\{\uparrow|\downarrow\}(\{I, R\}, \{I, R, T\}, \{+, -, ?\})$ representing the physical actions talked about in an encounter. \uparrow and \downarrow stand for a positive or negative link modification (i.e. addition/deletion of a link or an increase/decrease of its rating value), I/R for the initiator and responder of the encounter, T for a third party; $+/-/?$ indicates whether the (learning) agent likes, dislikes or doesn't know the target site of the link modification. For example, if a_1 and a_2 talk about `do($a_1, deleteLink(a_1, a_3)$)` in an encounter initiated by a_1 (while the learning agent a_2 is the responder and likes a_3 's site) this is abstracted to $\downarrow(I, T, +)$. If in the same conversation a_2 suggests to modify his own link toward a_1 (whom he does not like) from a rating value of 1 to 3, the state (viz *subject*) of the encounter becomes $\{\downarrow(I, T, +), \uparrow(R, I, -)\}$. The intuition behind this state abstraction method is to capture, in a generalised form, the *goal* of the conversation that can currently be realised while at the same time reducing the state space to a reasonable size.

Figure 4 shows a comparison for a population of ten agents and a fixed profile of private ratings (preferences) towards other agents (both plots show the performance of the best and the worst agent in the group as well as the average utility over all agents). The constants “upper benchmark” and “lower benchmark” denote the quantities for politically correct and honest linkage as discussed above, respectively.

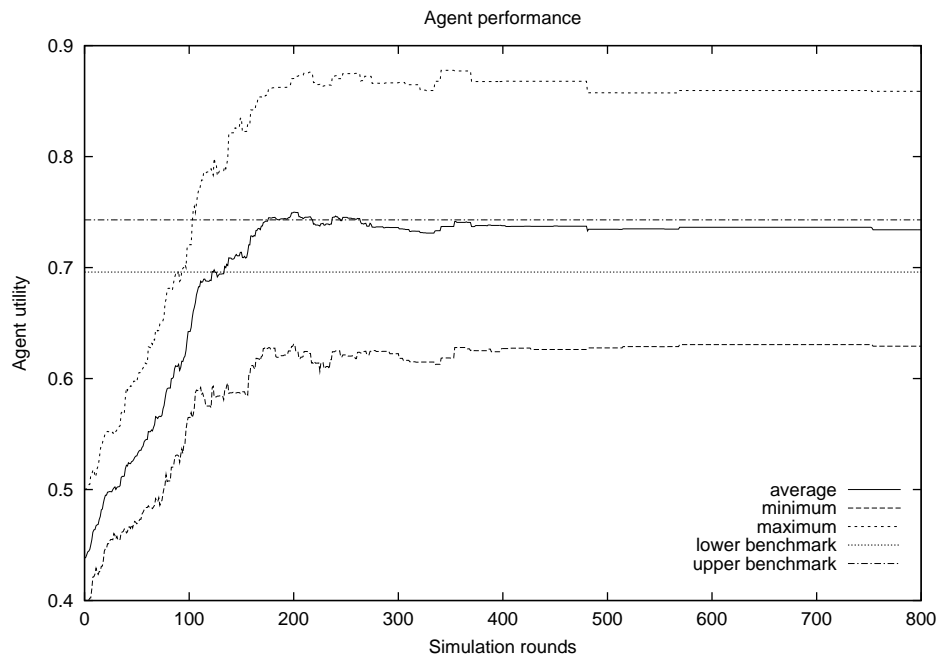



Fig. 4. Performance plots

In the first plot, agents employ BDI reasoning and additionally send requests to others whenever they favour execution of someone else's action according to their BDI queue. These requests are then enqueued by the recipient as if he had "thought of" executing the respective action himself. Thus, it depends on the recipient's goal queue and on his utility considerations whether the request will be honoured or not. As one can see, after a certain amount of time agents do no longer execute any of the actions requested by others, and cannot find any profitable action to execute themselves, either. The system converges to a stable state.

The second plot shows the results of a simulation with the same setup as above but using $m^2\text{InFFrA}$ agents. Again, agents issue requests whenever they identify that someone else could do something useful. Quite clearly, these agents perform much better than in the BDI case. In fact, more than one hundred simulations with identical settings have shown that we can ensure that agents will always reach a utility level close to that of the upper benchmark and this is quite impressive considering that $m^2\text{InFFrA}$ agents know nothing about cooperation potentials, others' preferences or the like. As in the BDI case, they prioritise their goals and actions only according to projected scores and contact others whenever they need someone else to perform an action that appears beneficial to them.

At first glance, it may seem only natural that $m^2\text{InFFrA}$ agents outperform BDI agents who use very simplistic means of communication. The important insight, however, lies in the fact that *any* rational agent design would perform poorly in this situation that is based on an pre-determined specification of communication semantics. What these simple BDI agents do is nothing but assess the usefulness of requests given that requestees always perform the actions others ask them to do. If this assessment suggests that others' actions are more desirable than those one can perform himself, it is only rational to prefer requests to physical actions. And yet, because of agents' egotistical stance, this results in a complete deadlock of the system where agents even fail to improve their situation using those actions they can perform themselves. Thus, implicitly, they are presupposing that everyone will adhere to a pre-specified semantics (namely, that all agents strictly obey the $\text{request}(A, B, X) \rightarrow \text{do}(B, X)$ pattern) but cannot respond to a situation in which this is not the case (for example by giving up and not issuing any more requests).

In the light of these results, we can state our main conjecture: In environments with (possibly non-benevolent) peer agents showing non-stationary behaviour, the performance of agents using prespecified communication protocols or assuming a fixed semantics of communication can become arbitrarily bad, while $m^2\text{InFFrA}$ includes the ability to learn what to expect from peers in a specific interaction situation.

Yet another interesting interpretation can be drawn from the fact that the average utility of $m^2\text{InFFrA}$ agents lies within the range of the two benchmark values. The fact that agent utilities evolve around these benchmarks indicates that they truly strive to make strategic communication moves and exploit the advantages of concealing certain beliefs.

7 Conclusions

In this paper we have presented an architecture for reasoning about communication patterns within the framework of empirical semantics. From a practical point of view, this contributes to the “agent side” of ACL research, as it offers methods to construct and implement agents that are able to deal with given specifications of rules that govern the communication in a system. Three features are particularly interesting about our approach:

1. The $m^2InFFrA$ frames used in our formal model combine information about the surface structure of communication sequences, logical constraints (frame conditions) and empirical data (counters). This allows for great flexibility with respect to *what* is defined in commonly agreed protocols, semantics, etc. For example, in non-exploitable protocols as those used in mechanism design, it may suffice to just specify admissible message sequences (because it does not matter what the agents think, for example), while other types of interactions such as contracting may require commitment rules, descriptions of agents’ mental states, etc.
2. Agents are capable of exploiting past communication experience and to combine this experience with specifications of communication patterns that have been provided *a priori*. Thus, they are able to “start out” with a predefined set of patterns and to test to which degree their peers comply with them. This offers a major advantage in open systems where adherence to protocols, truthfulness constraints, etc. cannot be guaranteed. At the same time, it paves the way for the use of *machine learning* methods that make minimal assumptions about others’ behaviour and simply accumulate communication knowledge as they go.
3. The architecture combines decision-theoretic (probabilistic) reasoning methods with symbolic communication. This constitutes an important contribution to the practical design of communicating agents, but also an initial step to explore the special character of communication as opposed to general action. In [20], for example, we have argued for trading off optimality against predictability in communication so that more reliable communication patterns can occur. Using decision theory as a firm foundation for rational reasoning can help formalising what agents (should) do in communication and whether and in which way this differs from the models used in, e.g. game theory and economics.

A major advantage of our approach is that it allows us to combine the decision-theoretic power of RL models with the knowledge-based aspects of symbolic agent communication, interaction protocols and ACL research in general. It is this aspect that makes rational action and learning possible for high-level agent architectures that employ logical reasoning.

Not all aspects of research on $m^2InFFrA$ have been covered in the present article. An extensive treatment of the additional components required to use frame-based learning as part of a complete agent architecture can be found in [3]. One of the most interesting aspects discussed there is a generalisation method for frame trajectories, which uses cluster validation techniques [7] on the (possibly fuzzy) clustering a set of frames induces on the space of possible message sequences. This endows agents with the capacity to create frames for encounters not matching any existing frame and to extend

the use of these frames to similar encounters in the future by abstracting from individual instances. Details of the application of hierarchical reinforcement methods on $m^2InFFrA$ are covered in [19]. Also, we have tested the architecture using more complex argumentation-based negotiation frames in [18]. Building on the framework of *interest-based negotiation* [15], these frames enable agents to gather information about other's goals and assumptions, to argue about the proposals they make, to attack others' claims, etc. To our knowledge, this work constitutes the first attempt to use learning in multiagent argumentation, and it can be considered a major advantage of our framework that it allows for the application of machine learning techniques to such complex forms of communication as argumentation-based negotiation. Finally, in [9] we discuss the potential of translating general expectation networks to $m^2InFFrA$ frames and vice versa.

In the future, we intend to look at more complex models of interaction frames with trajectory models that are more expressive than simple linear message pattern sequences. Also, we want to investigate to which degree $m^2InFFrA$ agents are capable of exchanging meta-frame information to reach consensus about which frames to use. As far as learning capabilities are concerned, we would like to introduce methods for inductive learning of context constraints. I.e., rather than just generalise over existing frames in terms of trajectory surface structure and updating frequency counters and substitution lists, we would like to enable agents to generalise from the belief states in which frames where applicable (positive samples) or not (negative samples). Finally, a lot of work needs to be done on automatically transforming ACL and interaction protocol specifications into readily usable $m^2InFFrA$ frames. In particular, the $m^2InFFrA$ formalism needs to be adapted to established frameworks for ACL and protocol semantics specification to facilitate automated transformation methods.

References

1. J. L. Austin. *How to do things with Words*. Clarendon Press, 1962.
2. J. Farrell and M. Rabin. Cheap talk. *Journal of Economic Perspectives*, 10(3):103–118, 1996.
3. F. Fischer. Frame-Based Learning and Generalisation for Multiagent Communication. Diploma Thesis, Department of Informatics, Technical University of Munich, Munich, Germany, December 2003.
4. FIPA (Foundation for Intelligent Physical Agents). FIPA Communicative Act Library Specification, 2002. Available under URL <http://www.fipa.org/specs/fipa00029/fipa00037/SC00037J.pdf>.
5. FIPA (Foundation for Intelligent Physical Agents). Fipa Contract Net Interaction Protocol Specification, 2002. Available under URL <http://www.fipa.org/specs/fipa00029/SC00029H.pdf>.
6. Nicoletta Fornara and Marco Colombetti. Operational specification of a commitment-based agent communication language. In Maria Gini, Toru Ishida, Cristiano Castelfranchi, and W. Lewis Johnson, editors, *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-02)*, pages 536–542, Bologna, Italy, July 15–19 2002. ACM Press.
7. Anil K. Jain and Richard C. Dubes. *Algorithms for clustering data*. Prentice-Hall, Upper Saddle River, NJ, 1988.

8. J. L. Kolodner. *Case-Based Reasoning*. Morgan Kaufmann, San Francisco, 1993.
9. M. Nickles and M. Rovatsos. Communication Systems: A Unified Model of Socially Intelligent Systems. In K. Fischer and M. Florian, editors, *Socionics: Its Contributions to the Scalability of Complex Social Systems*, Lecture Notes in Computer Science. Springer-Verlag, Berlin, Germany, 2004. To appear.
10. M. Nickles, M. Rovatsos, W. Brauer, and G. Weiss. Towards a Unified Model of Sociality in Multiagent Systems. In W. Dosch and R. Y. Lee, editors, *Proceedings of the ACIS Fourth International Conference on Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing (SNPD-03), October 16-18, 2003, Lübeck, Germany*. ACIS, 2003.
11. M. Nickles, M. Rovatsos, W. Brauer, and G. Weiss. Towards a Unified Model of Sociality in Multiagent Systems. *International Journal of Computer and Information Science*, 5(2):73–88, 2004.
12. M. Nickles, M. Rovatsos, and G. Weiss, editors. *Computational Autonomy. First International Workshop (AUTONOMY-2003), July 14, 2003, Melbourne, Australia*, volume 2969 of *Lecture Notes in Computer Science*, Berlin, Germany, 2004. Springer-Verlag.
13. M. Nickles, M. Rovatsos, and G. Weiss. Empirical-Rational Semantics of Agent Communication. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-04)*, New York, NY, 2004.
14. M. Nickles, M. Rovatsos, and G. Weiss. Formulating Agent Communication Semantics and Pragmatics as Behavioral Expectations. In *Proceedings of the AAMAS Workshop on Agent Communication (AC-2004)*, Lecture Notes in Artificial Intelligence, Berlin, Germany, to appear, 2004. Springer-Verlag.
15. I. Rahwan, L. Sonenberg, and F. Dignum. Towards Interest-Based Negotiation. In J. S. Rosenschein, T. Sandholm, M. Wooldridge, and M. Yokoo, editors, *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-03)*, Melbourne, Australia, 2003.
16. A. S. Rao and M. P. Georgeff. BDI agents: From theory to practice. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95)*, pages 312–319, 1995.
17. M. Rovatsos. *LIESON – User’s Manual and Developer’s Guide*. <http://www7.in.tum.de/~rovatsos/lieson/users-manual.pdf>, 2002–2004.
18. M. Rovatsos. *Computational Interaction Frames*. PhD thesis, Department of Informatics, Technical University of Munich, under review, 2004.
19. M. Rovatsos and F. Fischer. Hierarchical Reinforcement Learning for Communicating Agents. In *Proceedings of the Second European Workshop on Multiagent Systems (EUMAS-2004)*, Barcelona, Spain, December 15-17, to appear, 2004.
20. M. Rovatsos, M. Nickles, and G. Weiß. Interaction is Meaning: A New Model for Communication in Open Systems. In J. S. Rosenschein, T. Sandholm, M. Wooldridge, and M. Yokoo, editors, *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-03)*, Melbourne, Australia, 2003.
21. M. Rovatsos, G. Weiß, and M. Wolf. An Approach to the Analysis and Design of Multiagent Systems based on Interaction Frames. In Maria Gini, Toru Ishida, Cristiano Castelfranchi, and W. Lewis Johnson, editors, *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-02)*, Bologna, Italy, 2002. ACM Press.
22. M. P. Singh. A semantics for speech acts. *Annals of Mathematics and Artificial Intelligence*, 8(1–2):47–71, 1993.
23. M.P. Singh. A social semantics for agent communication languages. In *Proceedings of the IJCAI Workshop on Agent Communication Languages*, 2000.
24. R. S. Sutton, D. Precup, and S. Singh. Between MDPs and semi-MDPs: A Framework for Temporal Abstraction in Reinforcement Learning. *Artificial Intelligence*, 112:181–211, 1999.