Towards a Multiagent System for Competitive Website Ratings

Matthias Nickles
Institut für Informatik, Technische Universität München
80290 München, Germany, nickles@cs.tum.edu

Research report FKI-243-01

Abstract

Websites usually appear to the web surfer like isolated islands: Their self-description and their hyperlinks linking to other sites reflect solely the requirements, opinion and preferences of the respective site owner, there is no explicit relation to the information, service and product offers of competing sites, and the assessment of the site by other users is unknown as well. In contrast to usual recommendation systems, this paper outlines a multiagent system which derives multidimensional website ratings from the possibly conflicting opinions of interacting *rating agents* witch compete in the assertion of their individual ratings against others to provide a "socially enhanced" solution for this well-known problem. The rating agents act as representatives for individual web users, interest groups, other web sites, private or public organizations, and represent their ratings in an open discussion forum which is attached to the rated web site. This forum is continuously observed by a rating instance which computes rich, social weighted *general ratings* of different abstraction level from the forum communications. In contrast to the results of usual *majority voting* based recommendation systems, these general ratings take into account the social structures like norms and roles which emerge from the agents communication process. In addition to the presentation of general ratings to the web surfer, the ratings and the social structures can also be queried by the rating agents themselves to improve their rating capabilities.

Keywords: Multidimensional rating, Collaborative filtering, Semantic web, Multiagent systems, Agents, Socionics

1 Introduction

A common problem for every web surfer is the lack of reliable, objective content ratings of websites: If a valuing site description is available, in most cases it is either just a categorization of the site regarding some web directory (which classifies the site as belonging to a theme like sports or politics) or regarding some censorship criteria, or the rating is provided by the site owner himself, which

makes it as useful as any other kind of advertisement. In contrast, *recommendation systems* [4] based on the statistical evaluation of surfing behavior and/or the site content try to ascertain the "objective" value of the site. *Collaborative filtering* recommendation systems provide filter criteria for site classification, which classify the rated site in terms like "*appropriate / inappropriate*" or "*interesting / uninteresting*", based on the surfing behavior of a more or less homogenous group of users with a common interest profile (*user community*) and/or a *majority voting* process (a kind of opinion poll about the respective website) [11], which is carried out by the surfers or implicitly through an examination of other websites which interprets each foreign link to the site as a vote (a well-known example for this approach is Google's *PageRank* technology [18]). As a supplement or as a competing approach, *content-based filtering* recommendation systems try to analyze the content of web pages (usually by means of keyword counting) and compare the results with the interest profiles of the surfers [3]. The main drawback of such filtering systems is usually their limitation to one-dimensional ratings (amounting to "*like/dislike*") based on the presumed predilections of predefined or computationally defined social groups. This approach does not provide much help for the process of interest forming, which should in fact *precede* any filtering. Balanced ratings (i.e. regarding and weighting the opinions of multiple persons and groups) which are in addition reliable and unrestricted can currently only be provided by humans, like journalists, experts, or through discussion forums (e.g., newsgroups and threaded message boards like Slashdot (`www.slashdot.org`) or Everything$_2$ (`www.everything2.org`)), which makes this solution inappropriate for the huge and dynamic WWW, so that in most cases, there's no useful rating available at all. Another disadvantage of this kind of solution is the absence of a machine readable encoding of the results, making it almost impossible for internet information agents like web spiders to analyze the published ratings.

Although the W3C *semantic web* effort [17] addresses the problem of missing machine-understandability of web site descriptions, it currently focusses primarily on the specification of languages and tools for the representation of semantics and ontologies, not on the process of information gathering and valuing itself, and it doesn't care much about social phenomena like conflicting opinions, information biased by interest, and inconsistent or intentionally incorrect information.

In contrast to all of the described approaches, our goal is to provide so-called *general ratings*, that are *socially controversial* (i.e., achieved by the contribution of multiple, conflicting opinions represented by interacting *rating agents*) and *rich* (i.e. with unrestricted multidimensional rating criteria [2], multiple levels of generalization and unveiled information about the social relationship of their contributors), comparable with the résumé of a human discussion.

To achieve this, we make use of a new interdisciplinary field of research called *Socionics* [9] which is settled at the border of Distributed Artificial Intelligence (DAI) and sociology[1]. Socionics uses sociological theories for the construction

[1]The interdisciplinary computer scientific and sociological socionics effort has nothing to do with the psychological theory of the same name.

of highly complex (i.e. open, heterogeneous and large) multiagent systems to improve their scalability, adaptiveness and robustness. Particularly, we orient towards the sociological *Theory of Social Systems* ("systems theory") [7, 8], which enables us to describe multiagent systems (like a discussion forum of rating agents) on the *system level* [6] - in contrast to the traditional focussing of DAI on the mental properties of individual agents. This novel approach allows us to derive supra-individual general ratings which emerge from the communication of multiple subjective ratings. Accordingly, we will describe the participating agents only by means of some very general characteristics (especially autonomy), and instead focus on the facilitation of ratings through communications. This orientation towards the social system-level is also due to the fact that we want to describe an open MAS for heterogenous black-box agents, which are supposed to represent very different clients and which are equipped with very different capabilities and behavior. Thus it would be generally impossible to specify the rating agents in a uniform way regarding their "mental" properties.

This paper is structured as follows: The next section specifies the general requirements for a competitive ratings MAS. In section 3, we introduce the overall architecture of a framework for agent-based competitive ratings. Section 4 describes how social structures and general ratings can be founded in communication and interaction. Finally, section 4 outlines an agent communication language (ACL) for competitive ratings (*CRML*). With a discussion of open problems and directions for further research in section 5 the paper closes.

2 Terms and requirements

Basically, the rating of a certain website[2] is a finite set of so-called *elementary ratings*[3] together with a rating *vocabulary*. Each elementary rating is a statement about a *resource* (which can be any kind of object as well as another statement, like a picture or a text, or even another rating) by means of *meta data* (i.e., data about data), according to the vocabulary of *property types* (the set of rating criteria, also called rating *system*, *schema* and *ontology*). Technically, such an elementary rating is defined as a proposition of the form (*resource*, *propertyType*, *value*), in which *resource* denotes an object with an unique *Uniform Resource Identifier* (URI)[4] - in our context either a whole website or a somehow affiliated sub-resource (for example a single HTML page or the person which created the site). *propertyType* is the rated aspect of the resource (an element of the given rating vocabulary), and *value* is the assigned value of the property. *value* can be a resource by itself, and thus the rating vocabulary can form a hierarchy of property types (e.g., *Author* and *Credibility* in (`www.somesite.com`, *Author*, `John`), (`John`, *Credibility*, *high*)[5]). For each

[2]Corresponding to the *RDF* term "website *description*"

[3]Corresponding to the *RDF* term "*statements*"

[4]For convenience, we'll give the URL or a plain name instead.

[5]In this example, only the second statement is a real, i.e. possibly controversial, rating, while the first statement is an ordinary resource description ("The author of this page is

elementary rating within a certain website rating, *resource* must either be the whole website, or there must exist another elementary rating which contains the respective resource as its *value* component. Elementary ratings can be expressed in a formal *description language*, for example the XML-based *Resource Description Framework* (RDF) [10] (a successor of the internet rating language *PICS* [16]) or *DAML* [15]. An example for a rating with boolean property type is (`www.disney.com`, $ChildOrientation$, $True$). Beside such simple statements, rating languages like RDF can express more complex resource descriptions, but the principle is always as described, so we don't deal with technically more complicated kinds of ratings.

As an extension of the elementary rating, a *competitive rating* is defined as an elementary rating which is actively supported by a *rating agent* (who doesn't necessarily believe the rating he represents) against ratings represented by competing agents, with the goal to make some supra-individual *general rating* consistent with his individual ratings (this process we call "competitive rating" too). A *general rating* is a (probably inconsistent) website rating together with a social weighing of each element, i.e. a more or less abstract description of the degree of social assent and of the supporters or opponents of the respective elementary rating. We do not speak of "common" ratings, because this would restrict us to overall consensual ratings, which are generally very unrealistic. A well known example for a general rating is the result of an opinion poll: "10% of the questioned individuals support opinion A, 90% support opinion B". The percentages are supposed to describe the social acceptance of each opinion by means of a majority voting, which also underlies most website recommendation systems, although these systems usually unveil only the "winner" of the poll. Additionally, usual recommendation systems are limited to the single property type *appropriateness* with regard to the presumed needs of certain web site consumers. While our approach can lead to similar looking results (if the abstraction level of the respective general rating is high enough), our ratings are derived from an active competition process among competitive ratings (i.e., their representing agents), not just from the evaluation of a vote, and the general ratings are multidimensional, allowing multiple rating criteria at the same time. We assume that there already exist social structures with a strong influence on the WWW, including interest groups, coalitions and conflicts, which are normally hidden to the web surfer. So our task is to make these latent structures explicit to web users and information gathering internet agents by means of a socially rich rating which describes not only if a site is "interesting" or "appropriate", but primarily *who says so* and *by which criteria*. In a larger scenario, combining the general ratings of a significant selection of web sites, this could amount to a "social map" of the internet beyond superficial technical facilities like hyper linkage, and could thus be an important contribution to the forming of the semantic web, going far beyond the classical approach of specifying overall obligatory ontologies and knowledge representation formalisms, turning towards the difficult area of indefinite, inconsistent, controversial *social*

John"), but we consider such descriptions to be elementary ratings too to allow the easy rating of sub-resources.

semantics. Although this paper cannot give instructions ready for implementation, it maybe can show some initial steps towards the outlined goals.

The characteristics of competitive ratings and the resulting general ratings can be summarized as follows:

- *Distributiveness and admittance of controversies*: Non-rating website descriptions are rarely inherently controversial, and they do in general not require distributed computation for social reasons. Ratings in comparison are subjective quality judgements with a high conflict potential. Competitive ratings facilitate such controversies, and general ratings show them up.

- *Pro-activeness and dynamics*: The competitive ratings which contribute to general ratings are no single "passive" statements like votes in a poll, but are instead continuously represented by deliberative agents which support them actively in a dynamic social process.

- *Complexity and hybridity*: Due to the size, the heterogeneity and the openness of the world wide web, the agent-supported description of web resources is a highly complex task. Rating even increases the complexity of the *information-rich environment* the web constitutes for the rating agent by means of human sociality which becomes visible through the spectrum of different agent clients and their controversial ratings.

- *Unveiling of social relations*: Even an individual rating does not only describe the rated resource, but also makes an implicit statement about the rating actor himself. The co-presence of competing competitive ratings strengthens this exhibition of the rating subjects (i.e., the clients of rating agents), because it allows the derivation of relationships among opponents and supporters by means of a comparison of the respective ratings and the evaluation of the argumentation processes. This is supposed to increase the unveiling of social structures among websites, web surfers, user communities, organizations and other internet participants significantly.

Starting out from these properties, a facilitation of competitive ratings has three crucial technical requirements:

1. Formal *rating languages* both for the (on its own) definite competitive ratings and the (presumingly indefinite) general ratings to allow the utterance of individual ratings by agents and the representation of unsure and inconsistent meta data, contributed by different opinion sources, respectively. The focus should be on opinion announcement rather than on cooperation and consensus finding like in usual agent communication languages, but the languages should also allow the information exchange between agents and should be compatible with current and future standards for the description of web resources, like RDF, OIL and DAML (cf. section 5).

2. Intelligent *rating agents* which are able to deliberatively rate websites in accordance with the mutable opinions, criteria and interests of their private/public/commercial... clients and to represent their individual ratings in a discussion with other rating agents.

3. A technical instance (*rating system*) for the facilitation of rating agents interaction and for the derivation of general ratings from these interaction.

In this paper, we mainly deal with requirements 3 and and 1. As a solution approach, we sketch a multiagent system that consists of rating agents which act as representatives for web users, user communities, other web sites (or their runners, respectively), or private and public organizations. Every rating agent supports a certain opinion about the web site and announces and defends it as a set of individual ratings in an open discussion *forum* which is assigned to the rated web site and provides the technical infrastructure for the agents interaction, thus constituting a multiagent system. Every forum is part of the *Rating System* (RS), a central software instance that contains a so-called *Social System Mirror* [6] (SSM), which observes the rating forums and continuously derives their social structures from the forum communication and general ratings of levelled abstraction and differentiation from social structures. Together with the rated website, these general ratings are presented to the web surfer (e.g. in a special HTML frame within the browser window or via some user agent), and to mobile information agents, for example the *web spiders* of internet search machines like Altavista or Google, as well. The rating agents also can obtain the current general ratings and in addition the current social structures. This knowledge is considered to be very important for the agents to let them intentionally avoid or achieve oppositional or conformist behavior in respect to other agents and social norms, and to find allies and opponents.

3 Architecture

Figure 1 shows the proposed architecture of a framework for competitive website ratings, which consists of rating agents, web sites and the central rating system.

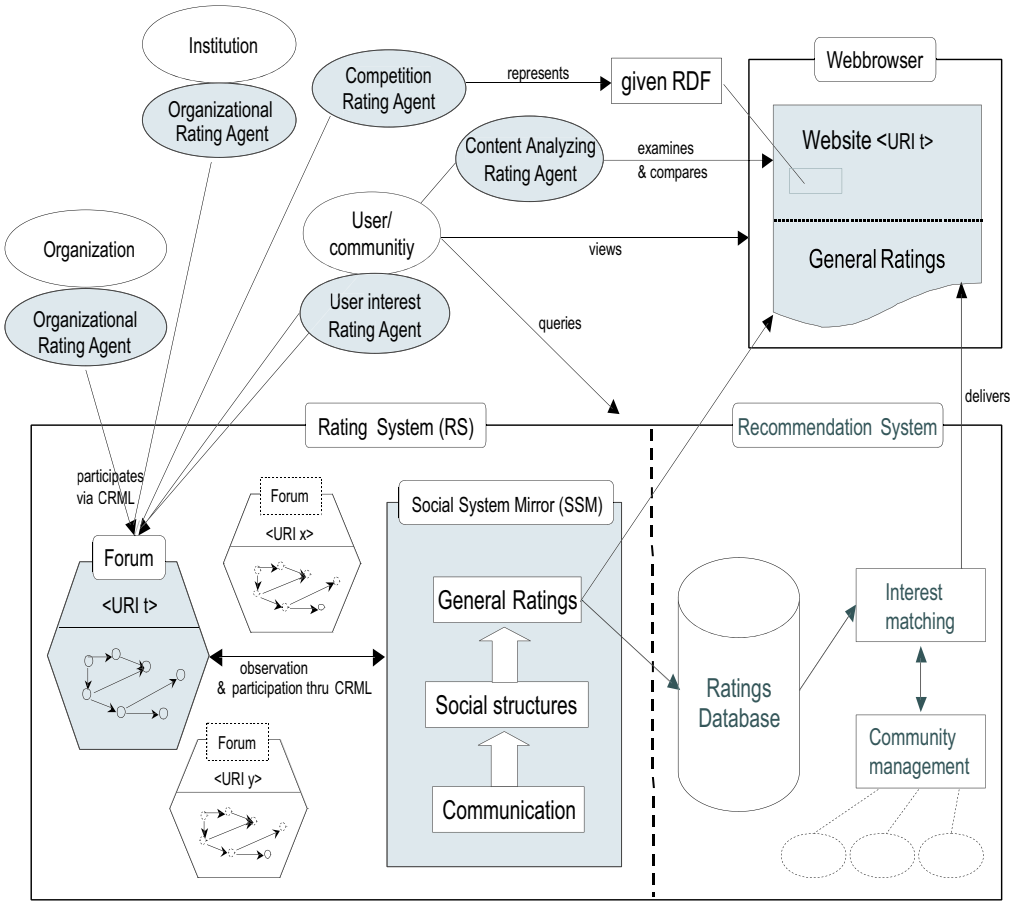

Figure 1: Proposed architecture of the rating framework.

3.1 Rating agents

In figure 1, the ovals in the top left part symbolize rating agents (gray) and their clients (white) - users, organizations and institutions (i.e., governmental organizations). Every rating agent must be able to form individual ratings conforming to the preferences of his client, to keep these ratings up-to-date with respect to changes of the website content and the opinion and needs of his client, and to represent the ratings in the rating forum. The latter means, that the agent "rationally" tries to influence the general ratings with the goal to make them consistent with his subjective ratings. This is only an indirect influence by means of inter-agent communication which is observed by the RS. To achieve his goal (which induces a so-called *goal conflict* with competing rating agents), the agent must have adequate technical, mental and social capabilities (for example the capability to cooperate with other agents in the assertion of some common rating). Since the general ratings are derived solely from the forum

communications, the agent must especially be able to participate in the forum (as a mobile agent which is executable on the forum server), master the rating communication language and know the algorithms the RS uses to derive general ratings to form a reasonable argumentation strategy. In addition, powerful rating agents not only have to be able to spread some predefined rating, but also to gather new information about the respective website, other comparable websites, useful information resources like internet databases, and in particular the ratings, knowledge and strategies of other rating agents. This is not only useful for them to form and improve the own rating and strategy, but also to inform the human or organizational client of the respective agent. Thus most agents are rather "two-way information agents" than just the mouthpiece of their clients[6].

The following set of partially overlapping major agent types is considered to be useful within a realistic world web environment. Of course, there will likely be an overlapping of these agent types in practice.

- *Organizational rating agents* represent private or public organizations (including companies) and their websites. They can represent interests and valuation criteria of almost any kind (commercial, legal, scientific, ethical, political...).

- *Organizational rating agents for institutions* are supposed to have a strong influence on other agents and the general ratings, because they represent governmental authorities, and their rating criteria and opinions are supposed to be based upon law. Such rating agents can participate in the RS to inform the web surfers (as far as the can influence the general ratings) and other agents about the legal assessment of the respective website, but also to label the rated site for e.g. child-protecting filtering software which relies on general ratings.

- *User rating agents* rate sites regarding the preferences of individual web surfers and communities of surfers with a similar interest profile. This can be done using a conventional rating by means of classical approaches of content- or collaborative-based recommendation systems, or alternatively in a combined process of simultaneous rating formation (e.g., through the evaluation of general ratings and observation of other agents) and rating announcement. It would also be conceivable to evaluate the content of the respective user home pages to derive their opinions and predilections. Like all rating agents, user rating agents are meant to act deliberatively, but can be implemented as simple "translators" of user intentions into formal rating language terms as well.

  Figure 1 also shows two kinds of agent which do not necessarily represent personal interests of human clients:

[6]We will not go deeper into these "mental" aspects in this paper, because we want to take an explicit system-level perspective which necessarily has to consider agents as more or less autonomous and heterogenous "black boxes".

- *Competition rating agents* compare the site with websites for which a rating is already available (which is of course not necessarily true or general). Two cases can be distinguished: If the site which is about to be rated and the comparative site are *identical*, the agent represents the known rating in the forum (as a kind of "advertising agent" for the site, if the existing rating has been attached by the site owner). Otherwise, the agent compares both sites and announces a weighted difference of the sites which expresses their relationship (like "This site is better/worse/... than the comparative site").

- *Content analyzing rating agents* examine the content of websites or subresources of them and feed the gathered information into the forum as asserted ratings. The information acquisition can be done via statistical keyword analysis (similar to the web spiders of web crawlers like Google, Altavista and Lycos), for example. Usually, content analyzing rating agents act as service agents that help other agents to develop or complete their ratings (but other kinds of rating agents shall be able to interchange such information too). Thus there are combinations of information rating agents and other rating agents conceivable, e.g. user interest rating agents which form a rating by means of a comparison of the results of a content analysis and the profile of their clients.

  Naturally, both competition rating agents and content analyzing rating agents are supposed to follow external hyper links on the rated website to find other sites for a further site comparison and examination.

3.2 The rating system

3.2.1 Concept

For each website, the rating system accommodates a forum which has the respective website as "theme of communication". The forum is technically implemented as a *whiteboard*, on which the agents write their messages addressed to other agents and read the responses - very much like people do in Usenet newsgroups and on web-based message boards. Every forum has its own rating vocabulary which has to be assembled by the agents themselves via some *ontology negotiation* technique see, e.g., [14]. Besides the syntax of the rating language and the communication protocol accompanying this language, the common vocabulary causes the only "hard" constraint for the agent's social behavior.

Forums are *open*, i.e. every agent who knows the rating language can participate in any forum at any time (even in multiple forums simultaneously), as long as he identifies himself and his clients to the RS and therefore to the system users[7]. As the central part of the RS, the Social System Mirror observes each forum, analyzes the communications and derives the social structures of the social system the forum (or multiple forums together) constitutes. In a

[7]For obvious reasons, the identity of the client should be authenticated by means of a trusted security certificate.

second step, for each forum (and therefore each website) the SSM generates general ratings from the social structures. As we will see in section 4.3, this is quite straightforward because general ratings can be considered to be certain subsets of the whole of the social structures. These derivations are carried out either continuously (similar to the examination and moderation of a newsgroup or a message board) or by request (when the respective web site is about to be displayed by the internet browser).

For the agent's communication a special *Competitive Ratings Meta Language* (CRML)[8] is proposed, which is outlined in section 5. For simplicity, this language is also used for the communication of rating agents with the SSM, i.e., to obtain information about social structures and general ratings. Since such communications shall not be visible for the other agents, they do not take place on the whiteboard.

The general ratings are not only announced to the rating agents, but primarily to the web surfer and to information gathering agents. Doing so, the RS and its SSM act very similar to a newspaper. Because of this, it is of course recommended to position the RS on "neutral ground", e.g. as a service of a web search engine which returns a general rating (encoded in HTML or XML) for each requested URL, or within a "rating proxy server" which delivers general ratings to the browser together with the requested web pages[9].
A possible extension of our framework which suggests itself is a recommendation system as part of the RS, which stores the general ratings generated from the SSM within a ratings database. Instead of asking for the general rating of a certain website, the user gives a personal interest profile to the recommendation system (e.g. in form of a desired rating). The recommendation system then compares this profile with each of the collected general ratings and returns a list of suitable URLs back to the surfer.

3.2.2 The Social System Mirror

The core of the RS, which derives social structures from communications and general ratings from social structures, is the *Social System Mirror* [6]. The "mirror" is a concept for a software instance which models the social system of the focussed multiagent system. It is not restricted to rating MAS, but can in principle be used with every system of autonomous agents. Technically, it can be thought as an intelligent knowledge base which derives social structures from communications and makes them available as information for the participating agents and other users. Originally a theoretical concept, it can be implemented as a software component (e.g. as a database system or as a middle agent). The mirror has two major tasks: First, the monitoring of agent communication processes and the continual derivation of emergent social structures from these observations, and second, the announcement of social structures to the agents

[8] CRML is called a "meta language" because it is intended to be implemented as an XML application which is compatible with the RDF.

[9] Instead of just publishing the ratings, it would of course be feasible to use them for filtering mechanisms which hide websites with a someway "bad rating" to the user, although this would obviously contrast the intention of our approach.

and the system users (the so-called *reflection* effect of the mirror)[10]. The agents can query the mirror very much like a database and deliberatively use the otherwise latent social structures which are made explicit to them through the mirror as a guide for their decision making and their interaction behavior, to assess their own social status and the status of other agents, to prevent (or intentionally not to prevent) the violation of social norms, and to selectively influence the current general ratings. If the agents deliberatively adopt to the reflected structures, those structures become stronger, otherwise weaker (the so-called process of *structure learning* in accordance with the evolution of social structures). By means of this the mirror reflects a realistic model of a social system to the agents and thus influences the agents[11], and the other way around, it continuously observes the actual multiagent system and adopts the social structures in its database in accordance with the agent's interactions. So the mirror not only provides emergent data about the global state of the MAS (like general ratings) which is supposed to be useful for the application user, but also improves the social skills of the participating agents. Doing so, it never restricts the autonomy of the agents. Its influence is solely through information and not through the exertion of pressure.

4 Social structures and general ratings

To derive general ratings from social structures like norms and roles, we need a formal description of "sociality" suitable for computation, a data structure for the representation of social structures, and algorithms to compute them from the interactions of the rating agents in the rating forum. The sociological *Theory of Social Systems* ("systems theory") of *Niklas Luhmann* [7, 8] provides an approach to these requirements because it understands sociality as constituted by *communications* (including social interaction in the wider sense), and treats social systems as *first-order objects* rather than just as a collage of mental properties of individual persons (or agents, respectively)[12]. Latter agent-centric "mentalistic" approach is still very common in DAI, but it is by no means appropriate for social agent systems with deliberative actors and emergent social phenomena [13], and it even leads to "profane" technical problems, if - as in our scenario - agents are more or less heterogenous and autonomous entities from the view-point of an external observer (like the RS) who cannot look "into their minds" - at least because of restrictions due to security and client privacy. An equally important reason for the usage of the systems theory for our application scenario is its universal modelling of social structures via *quantifiable action expectations*, which contrasts the usually more or less *ad hoc* descriptions of sociality in DAI. This allows us not only a relatively homogenous and simple representation of social structures, but also the direct transition from

[10]The announced structures are not necessarily the emergent structures derived from the system observation - the mirror can be pre-structured, i.e. used to reflect manually designed, non-emergent social structures as well, but this is not relevant in our context.

[11]Very similar to mass media in human society

[12]The computer scientific usage of this theory this text is based on is described more detailed in [6].

social structures (i.e., system-level action expectations) to general ratings (i.e., system-level rating expectations, cf. section 4.3).

4.1 Communication and expectation

It is widely accepted in DAI that the most important property of intelligent agents is their *autonomy* [5]. The major consequence of the autonomous behavior of agents is that a certain agent appears to other agents and observers more or less as a black box which cannot fully be predicted and controlled. This obscurity and uncontrollability is especially distinct in the open multiagent systems we are focussing on. Because only the actions of the agent in her environment are visible for an external observer, while her mental state is hidden to him, any *beliefs* and *demands* directed to the respective agent must be stylized as normative or adaptable *action expectations*, which are fulfilled or disappointed in future agent actions.

If it comes to an encounter of two or more agents, in which the agents observe each other, the described situation of mutual indeterminism is called *double contingency*. In trying to nevertheless determine the respective other agent and to achieve coordination (including "reasonable" conflicts), the agents have to *communicate*. A single communication is defined as the whole of a certain way of telling (a message act, which is not necessarily linguistic, but can also be a gesture for example), plus the communicated information, plus the understanding of the communication by means of a replying communication. Communication becomes visible for the observer as a sequence of agent interactions, that is, as a relation of message acts, and every kind of social interaction can vice versa be understood as a course of communications. Because communications are the only way to overcome the isolation of single agents, they basically constitute any kind of sociality, forming a *social system* of related communications (i.e. communications which refer to each other, like in a dialog). And if action expectations are related to message acts as parts of communications, the structures of a social system (consisting of communications) can be described as interrelated action expectations. That's why social structures are also called *expectation structures* if they are modelled as expectation interrelations [7].

Of course, communication (i.e. the interaction of agents) is not determined by expectations only, but expectations are used and adopted with every single communication because each communication makes a *proposal* (e.g., an assertion or a demand) which can be accepted or denied by the responding subsequent communication. Referring to expectations, this means that every communication relies upon existing expectations (including very basic assumptions about e.g. the common agent language) to form understandable and effective messages, and allows to test these expectations by comparing the observed subsequent message acts uttered by the other agent with the expected actions, and to form new expectations from this experience. This process of continuous structure adaption by means of interaction (or communication, respectively) is called *structure evolution* (cf. figure 4 in the next section).

Expectations regarding agent behavior can be formed not only as mental

states of other agents, but also from observers with a *global view* of the multiagent system. We call such *system-level* expectations *emergent* if they are solely formed from the empirical evaluation of the observed agent interactions (or communications, respectively). In our context, this global observation and derivation task is done by the SSM component of the RS for the purpose of constructing the social structures of the MAS.

4.2 Modelling of expectation structures

As we have seen, social systems like our rating forum consist of related communications, and the structures of the social system can be modelled as system-level expectation structures. In the following, we will outline selective aspects of an approach to the computational specification and representation of expectation structures and its application for the derivation of general ratings[13].

4.2.1 Elementary expectations

Basically, the social structures at *expectation time* are modelled as a set of simultaneous expected interdependent system-level expectations. Each expectation consists of an unique expected *action event*, the *actor* the expectation is directed to, and an unique triggering *situation* the event is expected to occur in. A tuple (*actor*, *situation*, *action*) is called *elementary expectation*. In our context, the action component is always the message act of a certain communication, which is usually the assertion of an elementary rating or the acceptance or denial of such an assertion. To denote expectations like "Agent X is expected to deny *all* requests" and "Agent X is expected to deny *some* request", it is allowed to place all- and existence-quantified placeholder within the action component. There are various ways to represent the situation part of an expectation, e.g. by time or by a complete or partial state description of the multiagent system and its environment. But since we derive expectation structures from communications (observable as courses of interactions), it is the obvious thing to describe situations as current or expected states of communication sequences (so-called *processes*). "Actual" situations correspond to certain preceding message acts the expected message acts are expected to reply to, "expected" situations correspond to message acts which are expected themselves recursively. So every expectation is *conditioned* by the assumption of a certain sequence of preceding action events (which is allowed to be incomplete or empty as well to express "a-priori expectations"), analogous to the conditioning of message acts by the preceding communication processes[14] (figure 2). If the expected action event occurs within a reasonable time-out span after the preceding message act and before alternative replies from the actor, the expectation is called *fulfilled*, otherwise *disappointed*. Since "physical agents" (i.e.,

[13] Because a detailed and formal description of this approach is beyond the scope of this paper, we have to refer the interested reader to [6] for a more detailed introduction to this developing field of research.

[14] This "conditioning" relation describes the representation of expectations, not their derivation. So the preceding process of an expected message act fixes the situation of expectation, but does not necessarily express the accumulated evidence for the expected event.

the software objects with their interior mental state) appear within social structures only indirectly by means of ascribed action expectations, we speak of a *social agent* being the actor component of expectations and denoting the set of expectations addressed to a certain physical rating agent. Single expectations directed to a social agent are called *elementary expectations*.

4.2.2 Spheres of communication and expectation

A set of communications and/or expected communications which are related by conditioning (i.e. if the partial order of conditioning shapes a graph with the respective message acts as nodes, like for X, Y and Z in "X has happened, and Y or Z is expected to happen subsequently") is called a *sphere of communication*[15] (e.g., the dialog of two agents). A sphere of communication which solely contains expected communications is called *sphere of expectation*. Such spheres are used to subdivide the current social structures for the purpose of a more subtly differentiated view on certain parts of the social system. For example, it is possible to derive and model expectations (especially addressed to social roles, see below) which are valid only within a certain sphere of communication.

4.2.3 Expectation strength, normativity and deviancy

Expectations can be weighted regarding their *strength* (that is, how much the action is expected) and regarding their *normativity*, or inversely, their *adaptivity*. Both the strength and the degree of normativity are real numbers ranging from 0 to 1. The degree of adaptivity describes how "ready to learn" the respective (generalized) expectation is if it was disappointed or fulfilled by some action event. Doing so, the normativity concerns primarily how much the expectation strength is increased (*reinforced*) or decreased (*weakened*) subsequently to the disappointment or fulfillment, respectively, but it can also affect the adaptivity of other potentially mutable attributes of the expectation, like the maintenance of the actor's social role. An expectation which is hardly adaptable is called *normative* (a "*norm*") and expresses how an agent *should* act rather than what he probably actually *will* do. So the strength of an adaptable expectation is simply the *a posteriori* probability of the expected action event (from the viewpoint of an objective external communications observer, in our case the SSM). In contrast, normative expectations are maintained in the longer term even if they are inconsistent with reality. So the strength of norms can be interpreted as a kind of "desired probability". In our case, the SSM establishes system-level norms not from its own demands directed to the agents, but instead tries to derive them from the observed communications[16]. The difference between the expectation strength of a normative expectation and the probability (i.e., the expectation strength if the degree of normativity would be zero) of a certain behavior is called *deviancy* and describes how much an actor deviates from the respective social norm (so, given the expectation strength and the deviancy, the

[15] Roughly corresponding to a *social sub-system* [7]

[16] Such communicated norms must not be confused with normative expectations of agents, because mental states do not necessarily become visible through communication!

real probability of a certain action can be calculated). There is a continuous transition from pure norms to pure adaptable expectations, and the degree of normativity of an expectation can change over time.
The representation of normativity and deviancy makes the crucial difference between social structures and other kinds of probability-based modelling of uncertain knowledge, like *Belief networks*. But apart from this, expectations can be calculated, represented and formalized analogous to probabilities. As an important consequence, the current set of expectations is not required to be consistent with the expecter's knowledge of reality, but it must be consistent by its own and valid in respect to the laws of probability. E.g. it is not possible to expect A and $\neg A$ at the same time for the same rating agent in the same situation if each expectation strength equals 1[17]. On the other hand, expectations which refer to expected propositions (i.e., propositions expected to be uttered by an agent) which contradict other expected propositions are permitted of course. Such indefiniteness can not only come from social conflicts: Because the calculation of expectation strengths is based on frequency analysis (cf. 4.2.4), an agent can make use of intentional inconsistencies (i.e., assertions with contradictious propositional content) to announce that she is uncertain about the proposition (the value of a rated property in our context). Another way to achieve this is to allow the agents to annotate their assertions with explicit "certainty factors", which are currently not provided in our expectation model.

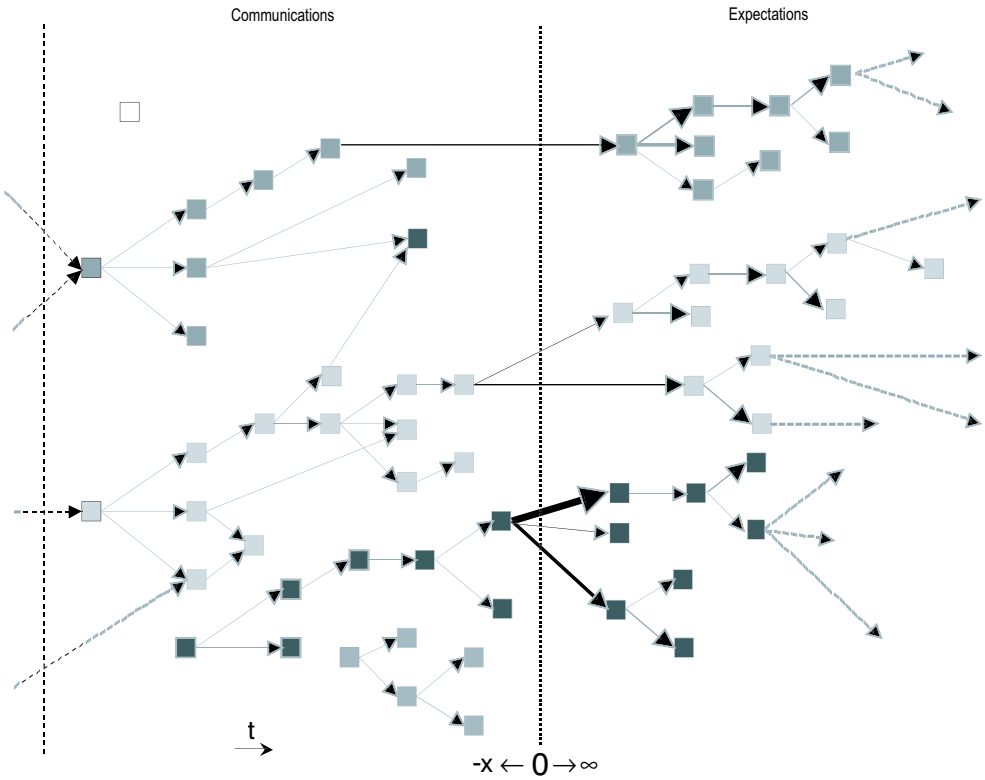

[17]The only exception is due to *social roles*, which can be limited to certain parts of the social system regarding their validity (cf. 4.2.4.1)

Figure 2: Communications and expectation structures at expectation time "0" in a simple graphical notion. Already occurred communication processes, starting at a certain initial time x (creation of the rating forum), consist of related message acts (symbolized as squares) and are shown on the left side of the dotted line in the middle. Expected processes in the future are shown on the right side. For the expectations, the thickness of the arrows connecting two message acts indicates the expectation strength. Complete spheres of communication are distinct by a different shading of the squares.

4.2.4 Actor and interaction course abstractions

In social systems, expectations are always directed to agents, but to gain a greater flexibility and independence, expectations can be represented in a more generalized style, abstracting from concrete agents and concrete courses of interaction. Beside the already introduced social agent, which is an abstraction of the "physical agent", systems theory distinguishes two other abstractions which are important in our context: *Social roles* abstract from social agents and *social programms* abstract from action courses. For both structure types, mappings exist to apply the assigned expectations to agents in concrete situations. Like elementary expectations, abstractions and expectations related to abstractions are derived by the SSM from the observation of agents interactions.

4.2.4.1 Social roles

Social roles ("roles") abstract from social agents. Every role ($roleIdentifier$, $agentsSet$) refers to a certain expected behavior (represented as a set of expectations with the respective role as the actor component or by means of participation in a social programm, cf. 4.2.4.2) which can be adopted by an element of the mutable agents set - the so-called *role instantiation*. While the mapping from physical agents to social agents is static and bijective, a role can not only be instantiated by multiple agents, but also every agent can impersonate multiple different roles (which even allows the social system to address otherwise inconsistent expectations to a certain agent) and the mapping from social agents to roles can change over time.
Like every expectation, expectations directed to roles (or to agents instantiating roles, respectively) (so-called *role expectations*), can range from a high degree of adaptivity to a high degree of normativity. In our model, the normativity of role expectations also affect the existence of the role itself (including the adaptivity of the agents set of the role): If the behavior associated with the role is expected highly normative, the agents are kept "imprisoned" within their role in the longer term even if their behavior deviates from the role expectations. If the role behavior is normative, the difference between the actual agent behavior and the expected behavior due to its assigned role is a form of deviancy, otherwise the role is adapted through the agent's behavior (presupposed the influence of the respective agent is high enough to change his role, e.g., if only a few other agents contribute to the role establishment and adaption) or the agent is removed from the agents set of the role due to the high difference between

the agent and his former role. Like social agents, roles can also be characterized by the size of their scope (measured as the cardinality of the largest sphere of communication they are contained in) and their stability over time (which does not necessarily comes from normativity). E.g., in the rating forum, *social positions* (a certain type of role) like the role of an "opinion leader" are assumed to be maintained over a relatively long period of time.
To make context dependent roles possible, it is allowed to have multiple roles with the same role identifier but different agents sets at the same time within different spheres of communication, such that multiple independent dialogs can contain the same role (e.g., the roles "Praiser" and "Critic" of different opinions in different dialogs). This is especially important in the (usual) case that the behavior of a certain agent is not completely determined by a certain role (e.g. the expectation, that a "baker bakes" does not mean that every person with this profession shall bake all the time). Context dependance can also be established by means of a semantical relationship of role identifier and the information parts of the communications addressed by the role expectations. E.g., the role "Conservative" could stand for rating agents which act in a conservative manner, possibly representing a rather political right-wing client[18].
To allow quantifying expectations like "*Every* agent is expected to..." and "*Some* agent is expected to...", two special roles are provided: The role *EachResponder* denotes every agent which responds in a certain situation (if the situation component of the expectation is undefined, *EachResponder* denotes every social agent currently participating in the rating forum). In contrast, the special role *SomeResponder* directs expectations to at least one unspecified responding agent.

4.2.4.2 Social programs

Social programs (including social plans and strategies) abstract from the behavior of single roles and agents in that they are flexible interaction schemes for multiple interacting social roles and/or social agents. Social programs usually have a high degree of adaptivity. In our context, they are not important for the derivation of general ratings, but are considered to be useful for the information of the rating agents by means of mirror reflections (cf. 3.2.2). Figure 3 shows a very simple social program, represented in a graphical notion.

[18]The set of agents of a role is not necessarily a *social group*, because group formation requires a social relationship among the group members by means of interaction. So group establishment via roles requires roles which are restricted to a certain sphere of communication. An exception of this rule is the special role "Member of group x".

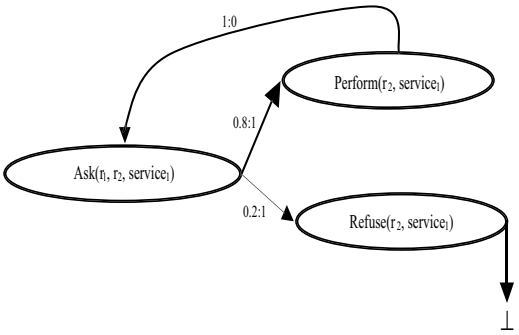

Figure 3: A simple social program

In this type of graph, the nodes correspond to message acts, which represent communications for their part. Message acts are modelled as speech acts, uttered and addressed by/to instances of roles (r_i) or agents. The directed arcs represent the expectation that the respective communication is followed by a certain subsequent communication. Arcs are weighted by expectation strengths (the thickness of the arc and the first labelling value). The second labelling value of each arc denotes the normativity of the expectation. Outgoing arrows at a certain node are always labelled with the same normativity, because their expectation strengths, analogous to probabilities, always sum up to 1 (if the expectation strength of one of the alternatives changes, the strengths of all other alternatives need to change by the same amount, that is, they have the same adaptivity). At $\perp$ the program terminates.

4.2.5 Derivation of expectation structures

Because the statistical derivation of social structures, which is carried out by the SSM, is far beyond the scope of this paper and many details are still object of further research, we will not go into the technical and statistical aspects of expectation derivation and just outline a possible approach.
Basically, the social structures (i.e., expectation structures as a set of emergent and related system-level expectations) of a certain social system like the rating forum can be calculated using a function applied to the set of all communications from the beginning of the system on until the present time, returning all expectations derivable from these communications. Due to the size and the complexity of social systems, this function can not be specified and applied in general. Instead, we have to use an *evolutionary approach* which continually updates an adequate subset of all social structures with every newly observed communicative agent message act in the rating forum. The main tasks of this mechanism are the establishment of new expectations, the adoption of current expectations by a recalculation of their strengths in accordance with their degree of normativity, the recognition and recalculation of the normativity itself, and the calculation of the deviancy of the actual behavior of the agents from their expected behavior. As described in 3.2.2, the updated social structures can be queried by the rating agents, and after each adoption step, general ratings are derived.

Expectations can be calculated as extrapolations of observed communica-

tion processes into the future, with the crucial assumption that similar communication processes will repeat themselves in the future (with an expectation strength proportional to the frequency of the already occurred processes, possibly in addition weighted by other communication attributes like actuality and regularity). Analogous, social roles and social programs are established through similar and frequent behavior and/or stereotype interaction schemes with multiple participating agents. To adopt expectations in dependence of the conformation or failure of previous expectations, they have to be *generalized*, i.e. the fulfillment or disappointment of an expectation regarding the current communication event must carry over to future situations also, which makes the evolutionary update procedure a kind of system-level *reinforcement structure learning*. E.g., the expectation that a certain agent will "deny all requests" generalizes over all future situations in which the respective agent is asked, and each actual denial or acceptance by this agent influences this generalized expectation.

For the practical calculation of generalized expectations, extended algorithms from linguistical speech recognition can be used, like the *n-gram extrapolation* [1], which have to be modified to predict dialog acts instead of phoneme or word n-grams. If the information part of communications is formalized in a computational language the expecter can master (first-order predicate logic), it is also conceivable to form and adopt expectations from the "inner" informational content of communications and not just from "outer" aspects like the frequency of message act sequences. E.g., if an agent agrees to the elementary rating (r, $ChildOrientation$, $False$) with expectation strength s, then it is reasonable to expect him to agree to (r, $ChildOrientation$, $True$) with strength $1 - s$. For an exhaustive unveiling of social structures it might be useful to further extend the capability to deduce expectation structures not only from the statistical evaluation of message act sequences, but also from semantical relations of the communicated information, as known from blackboard systems and deductive databases, but this would require an enhanced a-priori knowledge of the normative semantics of rating statements (e.g. wether a certain elementary rating is more general than another and therefore includes the less-general statement and makes is deducible this way).
The degree of normativity of a certain expectation can be quantified as the percentage of change of the generalized expectation if the expectation is disappointed. To recognize norms and to calculate the degree of normativity of expectations, sociological heuristics can be used (see [6] for examples of such heuristics).

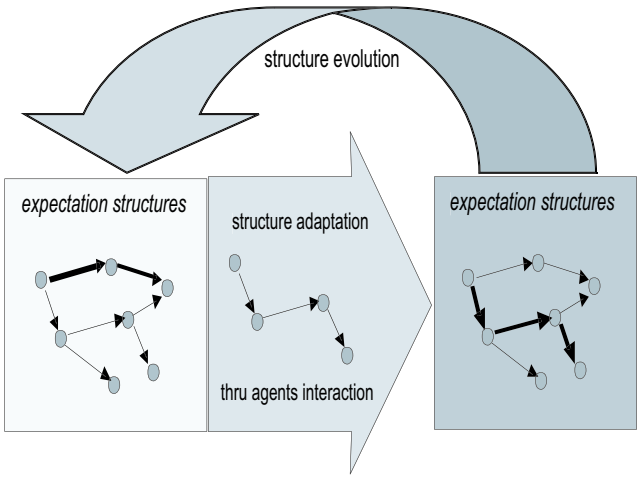

Figure 4: Evolution of expectation structures

4.3 General ratings

Now that we've identified social structures as expectation structures, it is very easy to derive general ratings from social structures. General ratings based on the underlying modelling of expectations outlined above are sets of *expectations regarding the utterance of elementary ratings* and therefore subsets of the social structures. Each element of a general rating is the system-level expectation of a *reply* to a question about the agreement with an elementary rating ($resource$, $propertyType$, $?value$), directed to an actor which currently participates in the rating forum (i.e., a social agent or a social role). To cover the opinion of the respective actor regarding $propertyType$, each general rating contains in principle all possible values of $propertyType$. Technically, this can be achieved via an enumeration of all possible instances of $?value$ for a given $propertyType$, determined by the rating vocabulary, and a corresponding query about the respective agreement for each instance.
The calculation of general ratings utilizes the fact that every multidimensional rating can be written as a set of one-dimensional *meta ratings*, i.e. "ratings of ratings" with a common property type component: Instead of (r, p, v) we can write ($p = v$, $agreement$, $degree\ of\ assent$). For example, the two instances of (`www.disney.com`, $ChildOrientation$, $?x$) are (`www.disney.com`, $ChildOrientation$, $True$) and (`www.disney.com`, $ChildOrientation$, $False$), and the expectation attributes of "positive" replies of an agent could be $strength = 0.9$, $normativity = 0.7$, $deviancy = -0.01$ in the first case (assertion of value $True$) and $strength = 0.1$, $normativity = 0.7$, $deviancy = +0.01$ in the second case (assertion of value $False$).

The following generalized expectations are supposed to be the basic elements of general ratings which can be calculated using our expectation model ($Accept$ denotes agreement, cf. section 5):

- ($agent_x$, $Query$: ($resource$, $propertyType$, $?value$), $Accept$)
- ($role_x$, $Query$: ($resource$, $propertyType$, $?value$), $Accept$)

- (*EachResponder*, *Query* : (*resource*, *propertyType*, *?value*), *Accept*) (i.e., a special role expectation which expresses the global acceptance of the respective rating.)

  All these requests can also be made restricted to certain spheres of communication, e.g.:

- (*EachResponder*, *Query* : (*resource*, *propertyType*, *?value*), *Accept*, $dialog_x$) calculates the overall acceptance of the proposition $propertyType \equiv ?value$ which emerges from $dialog_x$, under ignorance of the remainder of the forum communication.

Of course, a general rating shall usually not contain all ratings for all property types, and all agents, roles and communication spheres currently participating in the forum. Instead, it is reasonable to form general ratings of different levels of generalization. Although the ideal hierarchy of general ratings depends largely on the concrete application, the levels of meaningful general ratings could be within the range of the following extremes:

Top level (highest generalization):
actors≡*EachResponder*,
sphere≡whole forum,
property type≡the property with the highest overall expectation strength for its values

↕

Basic level (most details):
actors≡every single role and agent,
sphere≡every sphere with size larger than some ε,
property type≡every property type rated until now

To provide a machine-readable format for general ratings, it is reasonable to extend a traditional resource description language like RDF. This can basically be done by means of a replacement of the description parts of statements with *indefinite* propositions, which enumerate inconsistent elementary ratings as elements of a disjunctive set, and of the extension of values by expectation attributes. The rating vocabulary for the following example is implicitly given as XML namespace V, provided by some organization "description.org":

```
<rdf:Description about='http://www.disney.com'
xmlns:s='http://description.org/schema'>
     <V:ChildOrientation>
          <crml:disjunctive>
               <crml:boolean strength=0.9 normativity=0.7 deviancy=0
                    agent='Entertainment industry'>True</crml:boolean>
               <crml:boolean strength=0.1 normativity=0.7 deviancy=0
                    agent='Entertainment industry'>False</crml:boolean>
               <crml:boolean strength=0.3 normativity=0.7 deviancy=0.6
                    agent='User community 6'>True
```

```
                    </crml:boolean>
                <crml:boolean strength=0.7 normativity=0.7 deviancy=0.6
                    agent='User community 6'>False
                    </crml:boolean>
            </crml:disjunctive>
        </V:ChildOrientation>
</rdf:Description>
```

Another possibility is the usage of *statements about statements* (higher-order statements) provided by RDF. Here, we treat the elementary ratings as subjective statements, as in:

```
<rdf:Description>>
     <rdf:subject resource'http://www.disney.com'/>
     <rdf:predicate resource='http://description.org/schemaChildOrientation'/>
     <rdf:object>True</rdf:object>
     <rdf:type resource='http://w3.org/TR/1999/PR-rdf-syntax-19990105Statement'/>
     <a:attributedTo>Entertainment industry
         <crml:expectationAttr strength=0.9 normativity=0.7 deviancy=0/>
     </a:attributedTo>
</rdf:Description>
```

5 Outline of a competitive rating ACL

In the following, we sketch the basic characteristics of a proposed rating agent communication language CRML, which is supposed to be useable as a language for inter-agent communication *and* for the querying of social structures (in form of system-level expectation, with general ratings as a special case) by agents and the RS itself equally. The language is speech act - oriented and contains "social" speech acts which contribute to the derivation of social structures by the SSM as well as speech acts for the query of the SSM.

Basic social speech acts. $*$ is used within broadcast messages addressed to all agents:

Type	Receiver	Body	Subsequent	Illocution
Register	RS	*agent* identifier	-	Enter forum
Unregister	RS	-	-	Leave forum
RequestVocabulary	RS	-	-	Get property types[19]
Assert	*, *agent*, *role*	*proposition*	-, *Accept*, *Deny*	Assert proposition
Ask	*agent*, *role*, *	propos. *pattern*	-, *Assert*, *Deny*	Ask for completion
RequestSupport	*agent*, *role*, *	*proposition*	-, *Assert*, *Deny*	Request support
RequestDeny	*agent*, *role*, *	*proposition*	-, *Assert*, *Deny*	Request denial
Accept	*agent*	-	-, *Accept*, *Deny*	Deny previous
Deny	*agent*	-	-, *Accept*, *Deny*	Accept previous

The propositions within the bodies of the assertive speech acts are supposed to be encoded in a enhanced version of RDF (thus CRML and RDF have the same relationship as KQML and KIF). Since RDF is an XML sub-language (a

[19] Messages for negotiating the rating vocabulary are omitted.

so-called *XML application*), CRML messages should be encoded in XML also. In the table, "propositions" are always elementary ratings, i.e. statements of the form (*resource*, *propertyType*, *value*) as defined in section 2, or logical junctions of such ratings. In addition, there should be propositions ¬(*resource*, *propertyType*, *value*) for the utterance of negative assertions. For the *Ask* speech act, the "propositional pattern" can be derived from a rating proposition via the replacement of the value component with an existential quantified variable (which also requires an enhancement of RDF). Each speech act can be restricted to a certain sphere of communication, which is omitted within the tables due to lack of space.

Since we treat agents as black boxes, we cannot specify the semantics of the CRML in terms of illocution and perlocution. Instead, the semantics of a CRML speech act is (informally) defined as *the effect its utterance has on the expectation structures* (including general ratings). This goes far beyond the basic "finite state machine" semantics (given by the communication protocol sketched in the "Subsequent" column of the tables), and also prevents the rather intuition-based social semantics by means of *social commitments* suggested in e.g. [13].

Another important difference between CRML and other ACLs like KQML, FIPA-ACL and especially languages for argumentation and distributed theorem proving [12] is the focus on *opinion announcement* to the RS which is implicitly achieved with every agent ↔ agent message exchange since the individual goal of each rating agent is the influencing of the general ratings. So, even if a message is directed to another agent, giving the agent's opinion, another "implied receiver" is always the SSM. For this, facilities for argumentation should stay in the background (any individual rating is, as described in section 2, primarily subjective and cannot be "proven" to be true for other agents in general), while the differentiated announcement of subjective ratings and the agent's opinion about someone else's ratings should be in the center. This is achieved through the possibility for a certain agent to utter multiple (even mixed) assertions and denials referring to the same proposition to achieve a levelled degree of agreement, through speech acts like *RequestSupport* (which has a stronger supposed illocutionary power than *Assert* and therefore a stronger influence on the derived expectation strengths also), and the possibility for the agents to explicitly *Accept* and *Deny* any kind of statement and request - even another denial or acceptance. This allows a broad spectrum of conflicting behavior, rather than a broad spectrum of consensual or "consensus finding" behavior.

Because social structures structure communication processes (observable as subsequential message acts), it is reasonable to allow the query for such structures using the *same* language which is used for communications also. Analogous to the fulfillment or disappointment of expectations through the event of an utterance of one of the social speech acts in the first table, the speech acts within the following table can be used to request expected *future* communications subsequent to the current state of communication. These queries are useful for the agents as described in 3.2.2, but can also be used by the RS to query itself (respectively its SSM component) to derive general ratings from the current set of emergent system-level expectations. As allowed for social speech

acts, each query speech act can be restricted to a certain sphere of communication.

Elementary query messages[20]. The placeholder $*$ within this table can be replaced by each of the special roles *EachResponder* and *SomeResponder* (cf. 4.2.4.1):

Type	Receiver	Body	Result
QueryAgents	SSM	*role*, *	Agent in a certain role, all agents
QueryRoles	SSM	*agent*, *	Roles of a certain agent, all roles
QueryExpectation	SSM	message sequence *pattern*	Strength, normativity
QueryExepected	SSM	prediction sequence length	Message sequence
QueryDeviancy	SSM	time span	Deviancy
QueryDifference	SSM	two roles or agents	Difference
QueryRating	SSM	level	General rating

QueryExpectation(m_0 | *, m_1, [*,] m_2, [*,]...) calculates the strength of the expectation that a certain communication process will occur in the future (* stands for any sequence of messages). Every m_i has the form (*agent*|*role*|*, *message*), in which *message* can contain wildcards. m_0 has already to be occurred. Every m_{i+1} must be a response to m_i if it follows directly to m_i within the pattern. Example: *QueryExpectation*((*, (*,*Assert*(*p*)), (*Agent*$_1$,*Deny*))) asks for the system-level expectation that *Agent*$_1$ has the opinion that the rating p is wrong. *QueryExpectation* also returns the degree of normativity of the calculated expectation.

QueryExpected((m_0 | *, n) returns a sequence (*message*$_1$, ..., *message*$_n$) with the most expected communication process subsequent to m_0 respectively occurring "sometimes" (*). This type of message act can be used to obtain basic information about social programs, although more powerful query types are supposed to be required for this.

QueryDeviancy(*agent*, n | *) calculates the current deviancy (difference between expected and actual behavior for the last n resp. all messages.

QueryDifference(*agent* | *role*, *agent* | *role* | *) calculates the average strength difference between two sequences of expected behavior.

QueryRating(*level*) requests a general rating. This is just an abbreviation of the requests described in 4.3.

6 Conclusion

In this paper we have outlined a framework for a social rating MAS which is supposed to contribute to the emerging semantic web and to provide a novel approach to collaborative website recommendation. In contrast to usual approaches, we've focussed on the unveiling of social structures which arise from

[20] As in the table above, these are just examples to clarify the principle approach.

the communications of rating contributors and on detailed multidimensional general ratings instead of trivial filtering criteria. For this, we've defined social structures as expectation structures emerging from communication and have shown how rich general ratings can be derived from social structures.

This paper is by no means complete. A lot of work will be necessary for an exhaustive workout of the sketched tasks. Here just some of the open issues which are considered to be most relevant for future research:

- Many aspects, especially the rating language CRML and the derivation of expectation structures need to be formalized, enhanced and further specified.
- Currently, only small interaction systems (forums) are focused. On the way to the proposed "social map" of the WWW, a scaling of these local systems up to larger social systems (towards *internet society*) is required. It is uncertain, whether the naive merger of multiple forums (or websites, respectively) is sufficient to achieve this, or if further architectural enhancements of the framework are necessary.
- To facilitate complex communication processes beyond the announcement of subjective ratings, the CRML should be enhanced with language constructs for so-called *symbolically generalized communication media* [7] like "Money" (to allow utterances like "If you give me ..., then I will rate ...").
- Because the model of sociality underlying our focused framework is based on sociological theory, and due to the fact that many problems concerning competitive website ratings are inherently sociological and empirical, further sociological support of our computer scientific research will be necessary.
- A prototypical implementation of significant parts of our framework is required as a basis for a further elaboration of certain "fuzzy" aspects, like the establishment of social roles and programs, because these aspects cannot be completely specified without experimental support.
- Accordingly, the types of social structures like norms, roles and social programs need to be concretized with regard to our specific application scenario (e.g., which concrete argumentation schemes can occur? Which social positions of rating agents are expectable?).

Acknowledgements. This work has been supported by DFG (German National Science Foundation) grant BR609/11-1 (*Sozionik* priority program).

References

[1] P. F. Brown et.al. Class-based n-gram models of natural language. S. 283-298 in Proceedings of the IBM Natural Language ITL, 1990.

[2] R. Burkey. Semantic Ratings and Heuristic Similarity for Collaborative Filtering. Department of Information and Computer Science, University of California, Irvine, 2000.

[3] J. Delgado, N. Ishii, T. Ura. Content-based Collaborative Information Filtering: Actively Learning to Classify and Recommend Documents. In M. Klusch, G. Weiss (Eds.). Cooperative Information Agents II. Learning, Mobility and Electronic Commerce for Information Discovery on the Internet. Lecture Notes in Artificial Intelligence Series No. 1435. Springer-Verlag, 1998.

[4] J. Delgado. Agent-based Recommender Systems and Information Filtering on the Internet. PhD. Thesis. Nagoya Institute of Technology, 2000.

[5] S. Franklin and A. Graesser. Is it an agent, or just a program?: A taxonomy for autonomous agents. In J. P. Mueller, M. J. Wooldridge, N. R. Jennings (Eds.). Springer Verlag, 1997.

[6] K.F. Lorentzen, M. Nickles. Ordnung aus Chaos - Prolegomena zu einer Luhmann'schen Modellierung deentropisierender Strukturbildung in Multiagentensystemen. In K. Junge, T. Kron, S. Papendick (Eds.): Luhmann modelliert. Ansätze zur Simulation von Kommunikationssystemen. Leske & Budrich, 2001 (to appear).

[7] N. Luhmann. Soziale Systeme. Suhrkamp, 1989.

[8] N. Luhmann. Die Wissenschaft der Gesellschaft. Suhrkamp, 1998.

[9] T. Malsch (Ed.). Sozionik. Soziologische Ansichten über künstliche Intelligenz. Rainer Bohn Verlag / edition sigma, 1998.

[10] E. Miller. An Introduction to the Resource Description Framework. D-Lib Magazine, 1998.

[11] P. Resnick et.al. GroupLens: an open architecture for collaborative filtering of netnews. Proceedings of the conference on Computer supported cooperative work, 1994.

[12] M. Schröder. An efficient argumentation framework for negotiating autonomous agents. Proceedings of MAAMAW99. Springer-Verlag, 1999.

[13] M. P. Singh. A Social Semantics for Agent Communication Languages. In Proceedings of the IJCAI Workshop on Agent Communication Languages. Springer Verlag, 2000.

[14] T. Supnithi et.al. Learning Goal Ontology Supported by Learning Theories for Opportunistic Group Formation, AI-ED99, pp.67-74, 1999.

[15] http://www.daml.org/

[16] http://www.w3.org/PICS.

[17] http://www.w3.org/2000/01/sw/

[18] http://www.google.com