

Communication without Agents? From Agent-Oriented to Communication-Oriented Modeling

Thomas Malsch and Christoph Schlieder

Thomas Malsch, Technical University Hamburg-Harburg,
Department of Technology Assessment
Schwarzenbergstr. 95, 21071 Hamburg, Germany
malsch@tuhh.de

Christoph Schlieder, Bamberg University
Laboratory for Semantic Information Technology
christoph.schlieder@wiai.uni-bamberg.de

Abstract: From News to Chat, electronic discussion groups are widely acknowledged as a popular medium of communication. Unlike electronic mail which is rather easy to handle since it operates on a one-to-one bases, to keep up with forum discussions is extremely demanding. Participation in forums requires a constant effort of selection and attention from the user which goes beyond the limits of cognitive capacities. In this paper, we suggest to cope with this problem by introducing communication-oriented modeling (COM) as an alternative to agent-oriented modeling (AOM). Our approach to COM is based on theoretical foundations inspired by socionics and sociology.

1. Introduction

In this paper, we introduce a new approach into socionics and multi-agent systems research and design: *communication-oriented modeling* (COM). This methodological framework complements and reinforces *agent-based modeling* (AOM). In large-scale communication processes, especially those running on the Internet like discussion groups or chats, interaction between participants is often not organized along the lines of agent-to-agent relations. Rather, we find patterns of communication organized along the lines of message-to-message relations. Specifically, this can be observed in Internet-based public debates shaped by interrelated messages where a widely shared argumentation or a common view on a topic of general interest is gradually established (Lühns et al. 2001).

In Internet discussion groups messages usually are not sent to a specific receiver but "To Whom It May Concern". Messages are published to attract general attention and to enhance their social visibility by referring to other messages. Visibility in a general social sense, i.e. accessibility of a message and its potential of generating sequel messages, is a prerequisite of analyzing communication processes in real and artificial societies. Whenever a message is published for an audience rather than sent to a receiver, and whenever communication is dominated by messages referring to

other messages rather than agents influencing other agents' beliefs, intentions, and actions, it is communications rather than agents that should be considered as the foundational units of analysis and modeling.

The remainder of this paper is organized as follows: Section 2 gives an overview of practical problems arising in Internet discussion groups as seen from the perspective of participants and moderators. This is to illustrate the need to reinforce multi-agent platforms with methods and tools based on a new communication-oriented approach that does not depend on speech act theory along the lines of Austin, Searle and Habermas. In section 3 we introduce *communication-oriented modeling* (COM) as a methodological concept based on a social theory of evolving networks of communication which assumes that society consists of communication events rather than human beings. Section 4 elaborates the technicalities of COM in well known formalisms for logical and graphical representation. Section 5 gives an outlook on future research together with some hints on how to apply COM to multi-agent systems. Finally, to highlight the originality of our approach, section 6 gives an overview of related work in DAI, sociology and socionics.

2. Speech Acts and Agent-oriented Modeling

A prominent example for a communication process on the Internet – attractive to computer novices and experts alike – are the Usenet discussion groups which started in 1979, that is, more than a decade before the WWW. In December 2001, Google Inc. made its Usenet archive publicly available, thereby opening an incredibly rich source for studies in the history of computing and the sociology of communication. With more than 700.000.000 messages posted over a period of 20 years, the archive contains the best-documented large-scale communication process of the digital age. This makes it an ideal domain for illustrating some obvious limitations of *agent-oriented modeling* (AOM) of communication processes at a very large scale.

Agent-platforms such as FIPA-OS¹ (FIPA 1999) or JADE² are designed to enable communication between software entities (agents) showing goal-directed behavior rooted in a complex motivational system (BDI architectures). In the past few years, considerable effort has been made to develop agent communication languages (Singh 1998; Wooldridge et al. 1999; Bussmann et al 2000) in order to provide multi-agent systems with more transparency and coordinative power. The advantage of using agent-platforms for modeling communicative processes lies in the technical framework for agent communication which supports the purpose of communication analysis as well as that of simulation. An obvious way to model a Usenet discussion group within an agent platform consists in representing authors who are posting messages by agents of the platform. For modeling the messages, the unchallenged paradigm of communication in DAI is adopted: speech acts. Accordingly, along the lines of KQML standards (Finin et al 1994), agent communication is conceptualized as an

¹ FIPA-compliant open source platform distributed by Emorphia Inc. (fipa-os.sourceforge.net).

² FIPA-compliant open source platform distributed by TILAB Inc. (www.telecomitalia.it).

illocutionary act of a speaker (sender) who sends a message aiming at influencing the addressee's (receiver) intentions and actions. A similar sender-receiver pattern dominates computer communication in which the exchange of messages is regulated by protocols describing the precise conditions of starting the communication, acknowledging receipt, and so on.³ HTTP, the basic protocol of the WWW, for instance, adheres to the sender-receiver pattern as it regulates the flow of messages between a WWW client and a web server identified by an address, the URL (Comer, 2001).

In AOM, the minimal structure of a speech act can be described as being composed out of three components:

Table 1. The speech act in agent-oriented modeling

Agent1	Sender	persistent
Agent2	Addressee	persistent
Speech Act	Act(Proposition)	transient

At a more complex level, communication processes are specified by interaction protocols which are composed out of sequences of several speech acts. Thus, interaction protocols account for the fact that the addressee of a speech act is an autonomous agent too. It is realistically assumed that both, sender and receiver, are taking turns (quite in line with turn taking in conversation analysis). Platforms such as FIPA-OS are equipped with specific interaction protocols for different types of communication processes relevant in distributed problem solving (e.g. contract net protocol).

Although AOM is very successful with regard to distributed and cooperative problem solving, its speech-act based conceptualization of communication follows the sender-receiver pattern and, as a consequence, it has to struggle with a number of conceptual deficiencies and shortcomings when applied to large-scale communication processes beyond the scope of small-group interaction. In the following we will highlight three problems of the message sending paradigm underlying AOM.

1. *Focus on agent-agent relations.* In AOM, it is the agent who is considered to be the driving force of communication. The primary task of modeling consists in representing which agent authors a message (sender) and which agent interprets that message (receiver). Related design questions are: What is an agent's intention and how is it encoded in a message? However, in large-scale communication processes such as Usenet discussion groups, the intentional stance needs to be reinforced, if not substituted, with what we may call the referential or receptional stance: How is a message understood and how is it referred to by other messages? The necessity for this shift in focus away from agent-agent relations is supported by different empirical observations about discussion groups.

First, messages are not addressed to a specific receiver but are posted to be read by anybody who shows interest in them and invests the work for accessing them (message selection time plus interpretation time). This is in striking contrast to the mes-

³ A communication protocol "is specified by a data structure with the following five fields: sender, receiver(s), language in the protocol, encoding and decoding functions, actions to be taken by the receiver(s)" (Huhns & Stephens 1999, p. 86f).

sage sending paradigm. Second, and maybe less obvious, there is a tendency for the sender of a message to disappear in large-scale communication processes. Life and death of communications in a Usenet discussion group are largely independent from life and death of the individual agents participating in the discussion. Consider a typical Usenet group such as *alt.agnosticism* which started on July 1, 1998 and currently contains more than 69.000 threads with about two dozen new messages posted each day (significantly more on weekends). The independence of communication from individual agents is nicely illustrated by the fact that from the first 10 authors posting messages on the day the group started, not a single one has contributed to the discussion during the year 2002. The phenomenon of diminishing importance of the sender is also evidenced by senders which disappear behind pseudonyms or cryptic E-Mail addresses. Such participants could contribute to a discussion under more than one name/address, or a name/address might be shared by several participants.

2. *Missing message-message relations.* Another shortcoming of AOM and the message sending paradigm consists in not explicitly modeling the references that a message establishes towards other messages. The missing perspective of message-message relations causes the analysis of the communication process to take a specific turn. Agent-oriented analysis aims at describing agent-agent relations by structural or statistical means. A typical result would be a load distribution pattern in a communication graph whose nodes represent agents and whose edges correspond to messages having been exchanged between agents.

What cannot be extracted from the sender-receiver model of communication, however, is an explicit reference structure of a communication process relating messages to other messages. Because speech acts in DAI have not been introduced to refer to speech acts explicitly, message-to-message relations outside communication controlled by interaction protocols can only be established heuristically.⁴ Thus, a speech act primitive like "reject" that has been sent from agent *B* to agent *A* may be interpreted as a response referring to a "propose" previously sent from *A* to *B*. In an encounter of only two agents taking turns respectively, communicative acts indeed refer to each other according to the sequential flow of messages. In case of more complex communicative settings, however, taking temporal sequences for referential linkages is highly implausible. Whenever an agent exchanges messages with many other agents, perhaps along different protocols, or two agents exchange large numbers of messages asynchronously, not to speak of discussion forums with many participants addressing each other concurrently, heuristic referencing can no longer be considered to be reliable.

3. *High modeling complexity.* The figures behind Usenet discussion groups – a total of 700.000.000 messages with an annual increase of currently about 150.000.000 messages – render the task of modeling communication processes to be a prime challenge from the empirical as well as from the technical point of view. Huge amounts of data must be handled, a task impossible without computational assistance. More important, the data supporting the model at a given level of detail must be available. This poses a problem for AOM since, in general, nothing specific is known about the

⁴ In DAI, this difficulty seems to support "a view of the space of agent's interaction as merely the space of communication, ... where interaction histories simply result from the chaotic interleaving of the observable behaviours of single agents" (Ciancarini et al. 2000, p 250).

cognitive, motivational, and emotional state of the author who posts a message. To put it bluntly, there is simply no data available for modeling discussion groups on a very large scale in terms of goal-directed communication behavior of individual agents. Even if such data were available, the technical challenge of simultaneously running a very large number of agents (> 10.000) remains unresolved. Such a requirement is far beyond the capabilities of present agent platform technology⁵.

To sum up, we have identified three major deficits of AOM with respect to the modeling of large-scale communication processes. Most important is the observation that the continuity and outcome of communication in discussion groups cannot be explained as being warranted by agents continuously participating throughout the entire process. In contrast to what normally would be expected from cooperation in multi-agent systems with persistent agents on one hand and transient entities called speech acts on the other, agents appear to be transient in forums and discussion groups, whereas messages appear to be persistently available while the discussion goes on. It is the continuous availability of messages rather than the continuous presence of identifiable agents, that keeps the communication process alive and shapes its outcomes.

3. The Alternative: Communication-oriented Modeling

There is no doubt that agent-oriented modeling and speech-act theory have their own merits. However, with regard to analyzing and modeling complex social processes and structures as networks of communication in discussion groups, it has been shown that AOM exposes certain shortcomings and deficiencies. Hence we suggest that a different approach should be adopted: *communication-oriented modeling* (COM). Instead of modeling agent-to-agent relations we focus on modeling message-to-message relations as a methodological alternative to AOM. In our approach it is no longer the agent, but the communication event which is taken as the unit of analysis and design. COM has its conceptual foundations in a theory of communication which will be outlined in this section. This theory is, in turn, inspired by ideas taken from sociotics (Malsch 2001), from social theories of symbolic interaction and pragmatist semiotics (Mead 1934; Morris 1964; Mill 1998) and from Luhmann's sociological theory of social systems (Luhmann 1984; Stichweh 2000). It is based on three fundamental distinctions: inception and reception, observability and unobservability, and persistence and transience.

3.1 Conceptual Distinctions of a Theory of Communication Networks

1. *Reception and inception*. In our theoretical approach, communication is conceived of as a social process of messages linking or coupling to one another by *referencing*. Sociologically speaking, referencing is composed of two basic communicative operations called *reception* (understanding a message) and *inception* (producing a mes-

⁵ On a platform such as JADE some hundreds of complex agent can run simultaneously, with the widely used FIFA-OS this number is one order of magnitude smaller.

sage). Whenever a message visibly refers to another one, this is invisibly enacted by two subsequent operations: a predecessor message is received or understood and a successor message is received or produced. More formally speaking, we propose to define the term referencing as a temporal event: the moment when an edge is installed between two nodes, where the nodes are two messages while the edge is made up from a pair of two complementary operations, namely reception and inception. Thus, in our theoretical approach to COM and to social webs of communication, the unit of analysis and design is composed of two messages and two communicative operations.

Table 2. Unit of analysis in communication-oriented modeling

Message1	Message Sign	persistent, observable
Message2	Message Sign	persistent, observable
Referencing	Reception, Inception	transient, unobservable

2. *Observability and unobservability.* Any message can be seen from two perspectives: as a physical representation of an inception or as physical representation of a reception. A message, according to our theory of communication, is an empirically perceivable object. However, it is an object of a very special kind, namely a sign-object. Being a communicative sign, it designates to non-physical, meaningful, invisible communicative operations. In a very general sense, any message (gesture, spoken word, written text, picture, icon) must be construed as the empirical manifestation of unperceivable communicative operations. In contrast to empirically visible message signs, reception and inception are black boxed. Being operations that process meaning, they are unobservable. However, they can be reconstructed from relational constellations among message signs. Reception and inception must be distinguished by an observer who establishes a meaningful relationship between different message signs by referencing. Referencing is constituted as a meaningful relationship between sign-objects via reception and inception. In any given process of communication, receptions are linked to previous and successive inceptions, inceptions to previous and successive receptions, and so on, and as the process continues, a social network of communication is dynamically formed. It is the observer's or designer's task to open the black box and to explain how a communication network is enacted by its elementary operations.

3. *Persistence and transience.* Let us assume that communicative operations are transient. Communications come and go, one operation is followed by the next, inception is coupled to reception, reception to inception, and new messages are constantly added in a continuous process of social reproduction. Communications are elementary events, i.e. discrete, temporal elements in an ever evolving network of communication. Being events, communicative operations are temporally defined by the amount of completion time they need to process the meaning of a message. Being operational elements, inception and reception take exactly the amount of time they need to create or understand a message, e.g. read a book, utter a sentence, understand a question, write a letter. In an oral conversation, but also in Internet discussion groups, operations usually appear to be very short, ephemeral events. In a scientific discourse they tend to be much longer. In any case, they simply last as long as it takes to write or

read a paper. In both cases, however, in everyday life encounters of oral communication as well as in scientific discourses, any reception and any inception is processed discretely as a unique event which disappears when it is over. Of course, any operation may leave traces in memories and messages. But traces are traces of past events, not events in operation. In contrast to transient operations, messages are persistent objects, at least in textual and electronic communication. However, messages should not be misconceived as the immutable mobiles of social structure. Accordingly, a network of communication should neither be misconceived as having a static architecture. There is nothing static about social structures – in COM just as in real societies – since they are dynamically reproduced by new communicative events being added and old ones being deleted. A message's persistence or transience depends, in the first place, on its physical properties. What is more interesting, however, is that messages are activated, deactivated, or reactivated by selective referencing. This is what we call a message's relative social persistence or relative transience respectively. With regard to its social persistence, a message's social visibility can be enhanced and its life-span can be extended by being repeatedly referred to in subsequent messages. The more successor messages refer to a predecessor message, the higher its social relevance and significance, and hence its social visibility. And vice versa: any message may be socially deleted as a transient social object if it is constantly ignored or non-referenced, although it may continue to persist physically.

In the course of continuous referencing a social communication storage will be built up as an unintended (or emergent) infrastructure of communication. Unless it will be drawn on again and again by subsequent messages, sooner or later any given message or communication thread will be socially forgotten. Only think of Google's millions of inactive threads which have definitively disappeared from the social process of communication, although they are still "there". Note that a social network's survival does not depend on the survival of a single messages, however. As follows from the transient character of its operations and the socially constructed persistence of its messages, a social network's continuous reproduction also depends on the disappearance of operations and related messages, on selective referencing. A social network of communication, its structural persistence (and evolution) depends on the transient character of its operational elements. These have to be permanently activated in order to produce and reproduce masses of message-objects on a large scale basis in order to provide the network with sufficient redundancy which serves, in turn, as a prerequisite for evolutionary selectivity. Operations and messages must be continuously replaced by new ones to keep the social network alive. From a functional perspective, social reproduction (as well as innovation) seems to be in need of a permanent influx of new messages, although many of these, if not most of them, will never be reactivated or re-referenced again. Hence, sociality is run like a self-referential process: Messages selectively refer to other messages and in doing so they do not only permanently reproduce the operational elements of communication (reception, inception) and their empirical equivalents (sign objects), but they shape society as a dynamic structure which is both capable of stabilization and change.

3.2 Reception + Inception = Referencing

Viewing a message as a double manifestation of reception and inception is supported by empirical evidence about navigation behavior in Internet discussion forums. Usually, what we can see in the browser, are messages referring (explicitly or implicitly) to other messages, but not human participants referring to other human participants. Communicative events are not presented in an agent-oriented manner but in a problem- or rather argumentation-oriented manner. The way in which message threads usually are visually presented seems to indicate that something like a "gestalt switch" is taking place: from agents (and cognitive processing) to messages (and communicative processing). This does not necessarily mean to radically abstract from any idea of agent or agency. However, to make progress in COM, the agent's profile and persistency must be deliberately transformed into a background feature, while the message has to be switched from background to figure. Thus, the agent is no longer presented as the predominant figure and principal attractor of theoretical attention and design strategies. Or to put it differently: Only the message level is visible or empirically accessible while the level of operations is black boxed. Operations like reception and inception must be hypothetically or theoretically disclosed.

In line with such a "gestalt switch" it should be possible to reduce the amount of cognitive assumptions needed for COM to a minimum and to draw attention from psychological or cognitive processes to designing communication sociologically. It can be observed that most participants in Usenet discussion groups are actually silent most of the time. At the moment of posting a new message, a participant may explicitly or implicitly refer to one or more previous messages. In posting his or her message, he or she connects two distinctive communicative operations: a reception (of a predecessor message) with an inception (of a successor message). However, note that receptions do not automatically trigger inceptions within an agent's cognitive apparatus. In the contrary, any forward connection or junction from a reception to an inception is a highly contingent event if we assume agent autonomy. A message may be received, but the receiver may not be inclined to receive a new message. This may happen at any point in time. As a matter of fact, any encounter and any episode of interaction sooner or later comes to an end. As long as communication is faithfully recorded and physically stored, however, any reference structure for any given message can be reconstructed from past events, since any inception must have been triggered (incidentally or causally) by a previous reception. There is nothing like a first communication event.




Fig. 1. Ontological dependencies in communication-oriented modeling

To sum up our theoretical assumptions in a more stringent and formal way, we can say that, in Fig. 1, the arrows do not present causal relationships but stand for formal-ontological dependencies. These can be read in an upstream fashion or against the temporal flow of communication as follows: If X exists, then Y must exist. Whenever a message exists, it can be concluded that the inception which has created the message, must have existed too, and furthermore, also the previous receptions presupposed by the inception must have existed and, accordingly, the predecessor messages that have been received earlier. What we can observe here, is a curious gap of explanation between the inception of a message and its reception. The challenge is to precisely describe and explain the conditions under which inception is followed by reception. However, when we come to the next diagram about causal relationships in the temporal flow of communication, the theoretical problem seems to consist in the difficulty to explain why a reception is followed by an inception.




Fig. 2. Causal dependencies in communication-oriented modeling

To insist on the principle that every inception is based on a previous reception, admittedly implies to assume an actor who, for instance, gives a reply to a question. However, to point out to the fact that actors are always involved in communication and that the power to close the explanatory gap between reception and inception stems from actors and agency, misses the point in question. Viewing the issue from a functional perspective of large-scale processes of communication as in news groups, raises quite another question: What are the specific receptions an inception is drawing on and how can they be identified? Or, formulated on the level of perceivable messages: What are the particular messages from which a new message is generated or reproduced? Asking these questions means to proceed into a direction of analysis and design which is quite different from AOM. AOM is controlled by classical sender-receiver questions: Who sends what to whom? Or, in a derived form: Who are the favorite receivers of the messages coming from a sender? In contrast, to answer the questions of COM, we will have to explicate the reference structure of messages from underlying communicative operations or events. In doing so, it should be more convenient to leave the operational distinction between inception and reception behind, at least for the purposes of this paper, and to reduce both operations into a single operation: referencing. In the following section, referencing is introduced as the starting point of formalizing what we call message visibility.

4. Formalizing Message Visibility

According to our theory of communication outlined in the previous section, COM takes the relationship between a specific type of communicative events, namely the publication of messages, and the structure resulting from the references established between the messages, as its starting point. As we have just said, a publication event bundles the two complementary types of communicative operations which may be described as semantic actions that an autonomous agent is capable of performing: reception and inception. Instead of elaborating and formalizing the conceptual distinction reception/inception, it is more convenient for the purposes of this paper, to treat both operations as a single event of publishing or referencing at a higher level of abstraction. However, we will draw on the other two distinctions introduced in the previous section: observability/unobservability and persistence/transience. Hence, the formalization we propose mirrors the relationship between event and structure. It introduces two basic structures: the first describing the temporal ordering of publication events, and the second describing the reference structure of the messages. These descriptions lead us to a concept of *social* visibility which will be elaborated in this section.

4.1 Basic Structures of COM

The publication event structure (P, \leq) is a poset (partially ordered set) with the set of publication events P as ground set and the temporal ordering of events \leq as partial order relation. Intuitively, the partial order relation $p \leq q$ reflects the fact that p has been published before or at the same time as q . Note that the poset structure is equivalent to making the following assumptions about the temporal ordering of publication events:

Reflexivity	$p \leq p$ for all $p \in P$
Antisymmetry	$p \leq q$ and $q \leq p$ implies $p = q$ for all $p, q \in P$
Transitivity	$p \leq q$ and $q \leq r$ implies $p \leq r$ for all $p, q, r \in P$

In an application domain where a global synchronization mechanism exists which provides a unique time stamp for each publication event, the partial order becomes a linear order, that is, the following additional property holds:

Comparability	for any $p, q \in P$, either $p \leq q$ or $q \leq p$
---------------	--

The second structure central to COM is the *message reference structure* (M, \leftarrow) which consists of the ground set M of all published messages and a binary reference relation \leftarrow on M . Intuitively, $n \leftarrow m$ expresses that message m contains a reference to message n . Structural restrictions on the reference relation arise from the fact that messages are generated by publication events. This association is established by a bijection $\gamma: M \rightarrow P$ which maps a message m onto the publication event $p = \gamma(m)$ that generated it. Requiring γ to be bijective amounts to assume that there is no publication

event that does not produce a message (surjection), and that no two different messages are generated by the same publication event (injection):

$$\begin{array}{ll} \text{Surjection} & \text{for each } p \in P \text{ there is a } m \in M \text{ with } p = \gamma(m) \\ \text{Injection} & \gamma(m) = \gamma(n) \text{ implies } m = n \text{ for all } m, n \in M \end{array}$$

The references which a message establishes to other messages are restricted by the further requirement that they should be compatible with the temporal ordering of publication events. In other words, a message may refer only to messages that have already been published at the time of its publication.

$$\text{Compatibility} \quad n \leftarrow m \text{ implies } \gamma(n) \leq \gamma(m) \text{ for all } m, n \in M$$

Let us briefly discuss the implications of having messages inherit the temporal structure from publication events. For that purpose, we consider both, the publication event structure and the message reference structure, as digraphs (directed graphs) whose nodes are formed by the elements of the ground set and whose edges correspond to pairs of nodes linked by the temporal ordering or the reference relation respectively. Since the publication event structure (P, \leq) is a poset, the publication event graph does not contain any cycles. This property of being a DAG (directed acyclic graph) is inherited by the message reference graph because the bijection γ induces a subgraph isomorphism which embeds the message reference graph into the publication event graph. See Fig. 3 for an example of how the two graphs relate. In order to reduce visual complexity, not the publication event graph itself but its Hasse diagram⁶ has been depicted. Note that in general the message reference structure does not inherit the properties of reflexivity and transitivity from the publication event structure.

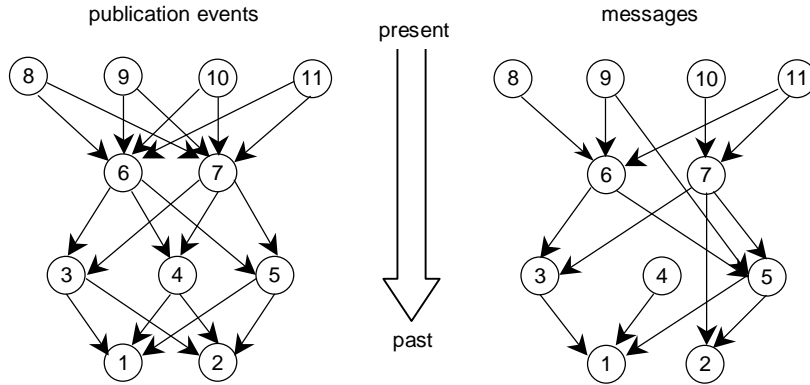


Fig. 3. Publication event graph (Hasse diagram) and message reference graph

⁶ To obtain the publication event graph from its Hasse diagram one must (1) add reflexive edges of the type $p \leq p$ at all nodes, and (2) add the transitive closure of all edges.

In view of the consequence of acyclicity one might raise the question whether compatibility is not too strong as requirement. After all, web sites frequently establish cyclical links between documents and scientific publications do list some of their references as being “in print”. With regard to this issue, we must clearly distinguish between the COM framework with its inherent assumption of acyclicity on the one side, and the way that this framework is used for modeling a particular application domain such as, for instance, web pages or scientific publications, on the other side. There is no doubt that in some domains, cyclical references or references to future publication events are useful or even necessary. However, from the perspective of COM such references appear as being composed out of a sequence of references which are acyclic and directed to the past only. If a scientific paper *A* cites another paper *B* as being “in print” then some preprint version *B'* of this paper must have been published before and known by the author of *A*. It is to this already published paper *B'* that the reference is established. The published form *B* of *B'* may now contain a discussion of related work with a cyclic reference to paper *A*. But this implies that the cyclic reference is constructed only in retrospective by an operation identifying *B* with *B'*. To sum up, we do not see empirical reasons to abandon the requirement of compatibility which is in agreement with our socionic theory of communication.

4.2 Measuring the Social Visibility of Messages

Communicative events (in our case: publishing events) do not persist over time although they leave a persistent trace in form of the messages they generate. A closer look reveals that the distinction is one of degree rather than principle: publishing does not occur instantaneously and messages do not exist forever in the sense that they do not remain eternally accessible for references from other messages. The COM approach claims that the empirical fact of temporally limited access to messages is not caused by the technical problem of making data objects physically persistent – an issue which is studied in the context of databases and digital libraries. Even with physically or technically persistent message signs, a decrease in accessibility will occur over time because the access to a message is linked to its social visibility in the communication process. The tendency of messages to become less visible over time is counterbalanced by the tendency of references to increase the social visibility of the message that is referred.

How exactly the temporal ordering of messages and their reference structure determine the visibility of a message depends on the specific application domain and the type of communicative process that is observed. We expect to find different measures of visibility for, say, Usenet discussion groups on agnosticism and on Java Server Pages, and – outside the Internet world – for scientific publications in computer science and sociology. For this reason we cannot but give one example among the many measures possible within our framework. It needs a few definitions which are likely to be relevant for other measures of social visibility too.

Let M be a message reference graph. We write S_m for the set of successors of a message m in M . In cases where the connection to sociocultural communication theory needs to be made explicit, we call S_m the *receptum* of m . The set of direct successors of m is denoted by DS_m . Analogously, we write P_m for the set of predecessors of m , also called the *inceptum* of m , and DP_m for the set of direct predecessors of m . Note that S_m and DS_m do not change over time whereas P_m as well as DP_m may be increased by new elements as new messages arrive which establish references to m . With most visibility measures, either increasing DP_m or increasing the inceptum P_m , i.e. the number of messages referring to m , results in an increased visibility of m . Note however, that also the receptum may play a role in determining visibility in some communication processes (e.g. citing outdated and exotic literature is not likely to increase scientific visibility).

Regarding the tendency for a decrease of visibility with time we only discuss the simplest case in which the publication events are ordered linearly. Furthermore, we assume that each publication event, and, as a consequence, each message m is associated with a real number $t(m) > 0$ that serves as its time stamp. Any monotonous decreasing function is a potential candidate for describing the decay of social visibility. A nearby choice for a function measuring the *recency of a message* is $r(m) = e^{-t(m)}$ which assumes the value 1 for the present, $t(m)=0$, and exponentially decreasing but always positive values for past events. This measure is easily integrated into a measure of visibility that also takes DP_m into account:

$$visibility(m) = \sum_{n \in \{m\} \cup DP_m} e^{-t(n)}$$

Note that this is just one out of many possible visibility functions which abstracts, for instance, from any influence that S_m could have on visibility. Fig. 4 describes the incremental construction of the message reference graph from Fig. 3. Four cycles of a simulation are shown. In cycle 1 two messages 1 and 2 are generated and assigned the visibility value 1.0 for new messages. The number of newly generated messages at a cycle c is a random variable $N(c)$ whose probability distribution is one of the parameters of the simulation model. We assume $N(c)$ to be equally distributed among values from the integer range $[1..4]$. In cycle 2, three more messages (3, 4, and 5) are generated. These messages establish references to the already existing messages (1, 2). The number of references that a newly generated message n establishes is again a random variable $R(n)$ whose probability distribution is another parameter of the simulation model. In this case, an equal distribution among values of an integer range $[1..3]$ was chosen. Visibility enters the play when it comes to determining the old messages that the references are directed to.

⁷ The definition uses the graph-theoretical notion of successor (and predecessor). Node 7 in graph M of Fig.3 has three successors, namely nodes 2, 3, and 5, as well as two predecessors, the nodes 10 and 11. This is not to be confused with the intended semantics of the reference relation. In an analysis of literary communication where the messages represent novels and the references are given by shared stylistic elements, novels 2, 3, and 5 may well be viewed as stylistic precursors of novel 7 although, graph-theoretically, they are successors.




Fig. 4. Visibility values during incremental construction of a message reference graph

Messages with high visibility are more likely to be referenced by new messages. The exact form of the distribution $V(o)$ describing the probability with which an old message o is referenced by a newly generated message is yet another parameter of the simulation model. However, the distribution must satisfy the following consistency condition for any two old messages o_1 and o_2 .

$$\text{Visibility} \quad \text{visibility}(o_1) \leq \text{visibility}(o_2) \text{ implies } V(o_1) \leq V(o_2)$$

In our case this is achieved by defining $V(o) := \text{visibility}(o)/\text{total-visibility}$ where *total-visibility* denotes the sum of the visibility values of all old messages. After the references have been established from the new messages (3, 4, and 5) to the old messages (1, 2), the visibility values are updated. According to the visibility function, newly generated messages are assigned the visibility value 1. For the old messages, visibility is diminished by temporal decay and increased by all incoming references. The original visibility of message 1, for instance, has decreased to $e^{-1} = 0.4$ to which adds increase of visibility by 3.0 due to the three incoming references from messages with visibility 1.0.

4.3 A Usenet Scenario

We close this section on the formal framework with an illustrative description of the constituents of a COM for a Usenet discussion process. Without detailed empirical analysis, it is, of course, not possible to come up with a fully fledged model that fits some set of empirical data. In particular, the probability distributions $N(c)$, $R(n)$ and $V(o)$ can only be determined with respect to a concrete communication process. The first step in COM consists in identifying the messages and the reference relation in the domain. A straightforward – but not the only possible – choice regarding messages is to adopt a realistic stance: each message posted in a discussion group appears as a *message* in the model. Similarly, the reference relation in the model represents, in the simplest modeling approach, *direct references* between messages in the discussion group. These are the references established by the author’s decision to post the message within a particular thread of a discussion group, generally as an answer to some other message.

Consider someone who wants to share a new type of argument refuting Creationism. This involves making a selection from an enormous range of possibilities. Among 100.000 Usenet discussion groups, the author opts for *alt.agnosticism* as the best place for publishing his argument. But this still leaves him with deciding which of the 69.000 threads to contribute to. He selects the thread “15 Answers to Creationist Nonsense” which contains more than 1.000 messages and decides to post his message as a comment to a message which already has produced some sequel messages, hoping that this will draw more attention to his message. Fig. 5 illustrates the tree structure of explicit references in Usenet discussion groups by showing the beginnings of the thread “15 Answers to Creationist Nonsense” in the discussion group *alt.agnosticism*.



Fig. 5 Threads in a Usenet discussion group

COM captures what is essential here: messages are published rather than sent. The crucial point is not that messages are made accessible to anyone who is interested but

that COM takes into account the fact that the amount of interest that any person can invest is highly limited, forcing everyone to make choices.

In a complex environment with thousands of messages arriving in a relatively short period of time, there is a limit to cognitive orientation. In other words, large-scale discussion processes exert a tremendous pressure on the cognitive capacities of human actors. Unlike electronic mail which is easy to handle since it operates on a one-to-one base, it is extremely demanding to keep up with a communication process which is inherently open and runs on parallel forums and threads. Participation in Usenet discussion groups requires an amount of selective attention from the user which far exceeds the limits of ordinary human capabilities. A user who wants to know what is discussed, whether there are parallel discussion groups on similar topics going on at the same time in different places, and when, where and how to post his own statement, is in permanent danger of getting lost. To compensate for limited cognitive capacities, COM could prove helpful since it suggests to reduce the complexity of Internet discussions in a way similar to what social networks do when constructing variable reference structures from an ongoing process of communication.

5. Future Work

As can be seen from the previous sections, there is still much work to do to develop COM as a powerful methodology and to demonstrate its scientific fruitfulness. In this section, we shall at first give an idea of how our approach to communication could beneficially be applied to current work in the multi-agent field. Secondly, we shall briefly address some more theoretical issues which seem to be particularly promising with regard to modeling and simulating artificial societies.

Interesting enough, not only an actor's cognitive orientation could be fostered by introducing methods of complexity reduction at the level of group interaction. What is more, strategies and techniques of complexity reduction and coherence management are of even higher relevance at the social level of large-scale communication. Take, for instance, the case of redundancy in parallel forums: Usually, it will be unfeasible for a moderator to prevent people from talking about the same issue in different threads or sub-forums at the same time. And even if so: Would an intervention be beneficial to the participants anyway? Who, in the role of a moderator, should be able to eliminate a posting as off-topic, or, in case of acceptance, who should decide where to insert which statement at which point in time or whether a contribution belongs to a specific topic or not? There is no global knowledge or privileged overview in a dynamic web of communication. At the global level there is nothing like a point of observation, whence everything could be seen as it really is. *Any observational perspective is part of the process and is, therefore, limited.* Consequently, "problem solving" in open forums like Usenet groups with mass participation cannot be viewed as being pre-structured by centralized coordination to such an extent that it leaves but some of the details to be settled by spontaneous contributions at the local level. In the contrary, the entire process is a contingent phenomenon. Any collective decision, any shared common view, consent or dissent, is but the outcome of communications referring to other communications. A software tool designed to assist participant users and

moderators has to be capable to observe and reconstruct a discussion in terms of messages referring to other messages. It should be able to register and monitor the flow of contributions, to analyze the process of new messages linking to previous ones, and to visualize temporal and referential patterns of communication. Additionally, an assistant tool should provide technical means of analyzing and simulating processes with very large numbers of communication events to study the effects of different scenarios on pattern formation and structuration. In the hands of a user, this should be an efficient instrument for individual orientation.

Moreover, in a technically more complex version, COM could be applied on top of a multi-agent platform. Consider a team of moderator-agents coordinating a large-scale debate in the Internet: How could these moderators profit from COM? What we suggest is to develop tools for distributed moderation. Obviously, moderating a discussion forum is more demanding than just preparing an individual statement and selecting the right moment for intervention. In order to come to grips with a communication process of ever growing complexity and redundancy, with different topics and opinions discussed in parallel, with multiple threads of argumentation, perhaps with sharp lines of conflict and dissent, there is a need for distributed coordination. To be successful, coordination should not try to control the debate but take advantage of *evolutionary tendencies towards differentiation or redifferentiation* as they occur in a quasi natural flow of communication. Hence, it might be helpful to distribute the moderator's task among several agents. Designing moderation as a cooperative task of coordination, begins with aggregating (or disaggregating) tasks and assigning each moderator agent with a specific responsibility. Instead of viewing the entire process with equal attention, one moderator agent might specialize on a specific feature, e.g. on conflict mediation, another focuses on a specific topic, a third one follows up a different topic, a fourth one analyses temporal patterns, turmoils and "hot spots" etc. To do so, all agents work with a COM-based communication analysis. However, they do so selectively and with different attentions. Again, this follows from the fact that *any observational perspective is limited*. Observing conflicts, for instance, means to distinguish between consent and dissent, irrespective of other features of the debate; temporality observation operates with the distinction immediate/postponed responses; topics observation distinguishes between on/off topic etc. Multi-agent moderation means that all these different observational distinctions will have to be cooperatively integrated, and this would be the task of an agent platform taking COM-based communication analysis as its input.

Coming to our theoretical perspectives, there are two issues that deserve particular attention in future work: social visibility, empirical observation, differentiation of sub-forums and reflexive or meta communication. These issues are closely interrelated and they just have been touched on in this section with regard to moderating discussion groups with COM and multi-agent systems.

1. *Visibility and observation*. A credible concept of social visibility should be based on the assumption that any observational perspective is limited and that there are always several perspectives installed. So far, we have considered social visibility as a global concept that applies to all publication events in the same way. This is also plausible for a small-scale process, e.g. a single thread in a discussion forum. In that case it can be assumed that all agents involved in publishing share – up to some toler-

able amount of error – the same measure of social visibility. However, in large-scale systems different perspectives of observation, hence different measures of visibility will coexist. In particular, the visibility of a single reference which forms the recursive base of the computation of the visibility function may differ. Consider a HTML document whose references, i.e. the hyperlinks, are visible only to software such as a web browser which can interpret HTML. The same holds for a human reader confronted with documents in different languages. Only some readers will be able to extract references from a text written in Japanese kanji characters. Therefore, an obvious refinement of our approach aims at allowing different publication events to work with different notions of social visibility. Another refinement concerns the computation of visibility itself. It is quite possible that forthcoming experience with modeling communication processes will show that other factors beside the temporal ordering of messages and their reference structure affect visibility. Obvious candidates are identifying operations on messages such as the one mentioned above which establishes the identity of content between different versions of a message. This is by no means a trivial problem. The fact that a message is part of a group of “identical” messages, e.g. different versions of the US constitution, contributes significantly to its visibility. In the same way, the effect of the author of a message on the message’s social visibility can be modeled as the effect of an identifying operation which groups messages by authors.

2. Differentiation and reflexive communication. We must try to develop models of social differentiation and re-differentiation to describe and simulate contrasting modes of dynamic reproduction, e.g. stable reproduction and evolutionary change of social structures in complex networks of communication. To achieve this goal, particular attention must be paid to the notion of temporality as expressed by the event structure of communication. With regard to social visibility values in the incremental construction of a message reference graph (Fig. 4), we can make the assumption that, for instance, a value lower than 0.4 means that a message falls under a given threshold of social visibility. Taking invisibility into account as an empirical phenomenon which is socially constructed in everyday communication by ignoring (non-referencing) a message, we need fine-grained patterns of referencing (adoption, rejection) recurring to the operational distinction between reception and inception again. From here, it should be possible to model and simulate social evolution: By selective referencing, a communication network gradually begins to separate into two (or more) different forums of communication, perhaps organized around different topics, opinions, authorities, authors etc. – different forums which will eventually have no longer any access to each other although they have been generated from the same home-forum. This means that we will have to study social invisibilization at a more specific level of investigation. Moreover, we need theoretical models of what sociologists would call the innate reflexivity of social communication. A discussion becomes reflexive when participants begin to communicate about the discussion as a communication process. Any communicated observation concerning the discussion at a meta level (rules, outcomes, topics, and standard) inevitably inserts reflexivity into the process. Obviously, reflexive communication may have an enormous impact on the entire process of communication, and therefore it must be regarded as one of the central tenets that we will have to cope with in future work.

6. Related Work

Our work is related to different fields of research in DAI, sociology, and socionics. To begin with socionics (Malsch 2001; Kron 2002), our approach is closely related to expectation-oriented analysis and design, based on modeling expectation by means of a social mirror (Brauer et al. 2002; Lorentzen & Nickles 2002). This work is, similar to ours, inspired by ideas taken from social systems theory (Luhmann 1984). In contrast to our proposal, however, the social mirror is still based on speech acts. Moreover, it does not explicitly allow to model the temporal structures of social change in terms of complex chains enacted between communication events.

Coming to DAI research, our approach has been encouraged by recent publications suggesting to put communication and interaction on top of the agenda for multi-agent systems design (Wooldridge et al. 1999; Bussmann et al. 2000). Our proposal differs from DAI's paradigm of communication dominated by KQML speech act primitives and FIPA standards. Communication in agent-oriented modeling has its focus on interoperability issues such as the common language problem (Labrou et al. 1999), conversation-type of interaction (Barbuceanu & Fox 1997), and dialogue-oriented communication in small groups of agents (Bretier & Sadek 1997). In our paper, we take interoperability for granted. Rather than dealing with interoperability issues, we want to contribute to the coordination problem of DAI. As has been recently suggested (Ciancarini et al. 2000), agent interaction should not be viewed as merely occurring within a given technical infrastructure of communication. In order to deal with the complexity of interagent communication more efficiently, theories and tools are needed to design coordination into multi-agent systems via social rules and collective commitments. In our paper we do not directly tackle the coordination problem by providing, for instance, agent ensembles with off-line designed social laws (Shoham & Tennenholtz 1995) or models of other agents' beliefs, abilities, and preferences (Gmytrasiewicz & Durfee 2001). We do so rather indirectly by studying a very simple on-line mechanism of a highly abstracted social structure that emerges as a temporal pattern of communicative events.

In fact, this is the central tenet of our paper, and it is clearly inspired by the sociological turn from action to social interaction and communication (Mead 1934, Luhmann 1984, Stichweh 2000). According to Luhmann's theory of autopoietic social systems, communication must be construed as the temporal element (or basic operational event) of social systems that reproduce themselves by permanently producing the very elements of communication they are made up of. Our proposal to represent social evolution as a combination of a reference structure and an event structure, is different in that it has been rather freely adopted from Mead's and Luhmann's views of society as a dynamically evolving network of communication. Moreover, our approach to COM has also been cross-fertilized with other sociological concepts of communication taken from conversation analysis, and objective hermeneutics (Garz & Kraimer 1994; Schneider 1994). Our distinction between perceptible messages and unperceptible albeit meaningful and hence accessible communicative operations is taken from there.

Last not least, our work is related to the study of Usenet discussion groups and other Internet forums with regard to democratic participation in large-scale public

debates (Lührs et al. 2001) and to methods of social network analysis (Wasserman et al. 1994) applied to internet discussions (Albrecht 2001; Lübcke 2000). A major deficiency of social network analysis (SNA) must be seen in the fact that it is based on an agent-oriented and rather static methodology. Indeed, SNA seems to suffer from a considerable lack of providing appropriate means to describe the evolutionary dynamics of social networks. Again, our own approach to COM is specifically designed to analyze and simulate evolutionary processes of network configuration and might, perhaps, contribute some ideas to render SNA more dynamic.

References

1. Albrecht, S.: Structuring Large-Scale Online Debates - Making Use of Network Analysis Methods. Paper presented at the Sunbelt XXI International Social Network Conference, Budapest, Hungary, April 25-28 (2001)
2. Barbuceanu, M., Fox, M.S.: The Design of a Coordination Language for Multi-Agent Systems. In: Müller, J.P., Wooldridge, M.J., Jennings, N.R. (eds.): Intelligent Agents III. Agent Theories, Architectures, and Languages. ECAI '96 Workshop (ATAL) Budapest, Hungary, August 1996 Proceedings., Springer-Verlag Berlin, Heidelberg (1997) p. 342-356
3. Brauer, W., Nickles, M., Rovatsos, M., Weiß, G., Lorentzen, K.F.: Expectation-Oriented Analysis and Design. In: Wooldridge, M.J., Weiß, G., Ciancarini, P. (eds.): Agent-Oriented Software Engineering II. Second International Workshop, AOSE 2001, Montreal, Canada, May 29, Lecture Notes in Computer Science, Vol. 2222. Springer-Verlag, Berlin Heidelberg New York (2002) p. 226-234
4. Bretier, Ph., Sadek, D.: A Rational Agent as the Kernel of a Cooperative Spoken Dialogue System: Implementing a Logical Theory of Interaction. In: Müller, J.P., Wooldridge, M.J., Jennings, N.R. (eds.): Intelligent Agents III. Agent Theories, Architectures, and Languages. ECAI '96 Workshop (ATAL) Budapest, Hungary, August 1996 Proceedings, Springer-Verlag Berlin, Heidelberg (1997) p. 189-204
5. Ciancarini, P., Omicini, A. and Zambonelli, F., "Multiagent System Engineering: The Coordination Viewpoint". In: Jennings, N.R.; Lespérance, Y. (eds.): Intelligent Agents VI: agent theories, architectures, and languages; proceedings / ATAL'99, Orlando, Florida, USA, July 15 - 17, 1999- Berlin; Heidelberg; New York: Springer, (2000) p. 250-259
6. Comer, D.: Computer Networks and Internets. Prentice-Hall: Upper Saddle River, NJ. (2001)
7. Finin, T. W., Fritzson, R., McKay, D. and McEntire R.. KQML as an agent communication language. In: Proceedings of the 3rd International Conference on Information and Knowledge Management (CIKM '94), Gaithersburg, Maryland, ACM Press (1994) p. 456-463
8. FIPA: Foundation for Intelligent Physical Agents. FIPA '99. <http://www.fipa.org>.
9. Garz, D. und Kraimer, K.: Die Welt als Text. Zum Projekt einer hermeneutisch-rekonstruktiven Sozialwissenschaft. In: Garz, D. und Kraimer, K.: Die Welt als Text. Theorie, Kritik und Praxis der objektiven Hermeneutik. Suhrkamp, Frankfurt am Main (1994) p. 7-22
10. Gmytrasiewicz, P.J. and Durfee, E.H.: Rational Communication in Multi Agent Environments In: Autonomous Agents and Multi-Agent Systems 4 (2001) p. 233-272
11. Huhns, M.N. and Stephens, L.M.: Multiagent Systems and Societies of Agents. In: Gerhard Weiss (ed.): Multiagent Systems. A Modern Approach to Distributed Artificial Intelligence, MIT (1999) p. 79-120

12. Kron, T.: Luhmann modelliert. Ansätze zur Simulation von Kommunikationssystemen. Leske + Budrich, Opladen (2002)
13. Labrou, Y. Finin, T. and Peng, Y.: The Interoperability Problem: Bringing together Mobile Agents and Agent Communication Languages. IEEE. In: Proceedings of the Hawaii International Conference On System Sciences, January 5-8, 1999, Maui, Hawaii (1999) <http://www.cs.umbc.edu/~finin/papers/hicss99.pdf>
14. Lorentzen, K.F., Nickles, M.: Ordnung aus Chaos - Prolegomena zu einer Luhmann'schen Modellierung deentropisierender Strukturbildung in Multiagentensystemen. In: Kron, T., (ed.): Luhmann modelliert. Ansätze zur Simulation von Kommunikationssystemen. Leske + Budrich, Opladen (2002) p. 55-113
15. Lübcke, M., Kommunikation im Netz. Eine Analyse am Beispiel von Newsgroups. Diplomarbeit im Fach Soziologie, vorgelegt bei der Universität Hamburg (2000)
16. Luehrs, R., Malsch, T., Voss, K.: Internet, Discourses and Democracy. In: Terano, T. et al. (eds.): New Frontiers in Artificial Intelligence. Joint JSAI 2001 Workshop Post-Proceedings. Lecture Notes in Computer Science, Vol. 2253. Springer-Verlag, Berlin Heidelberg New York (2001) p. 67-74
17. Luhmann, N.: Soziale Systeme. Grundriss einer allgemeinen Theorie. Suhrkamp, Frankfurt/Main (1984)
18. Malsch, T.: Naming the Unnamable: Socionics or the Sociological Turn of/to Distributed Artificial Intelligence. In: Autonomous Agents and Multi-Agent Systems 4 (2001) 155-186
19. Mead, G.H. (1934): Mind, Self and Society. Chicago (1934)
20. Mill, U.: Technik und Zeichen. Über semiotische Aktivität im technischen Kontext. Nomos, Baden-Baden (1998)
21. Morris, C.: Signification and Significance. A Study of the Relations of Signs and Values. MIT, Cambridge, Mass. (1964)
22. Schneider, W.L.: Intersubjektivität als kommunikative Konstruktion. In: Fuchs, P. and Göbel, A. (eds), Der Mensch - das Medium der Gesellschaft. Suhrkamp, Frankfurt am Main (1994) p. 189-238
23. Shoham, J. and Tennenholtz, M. Social laws for artificial agent societies: Off-line design. In: Artificial Intelligence 73 (1995), p. 231-252
24. Singh, M.P. Agent communication languages: Rethinking the principles. IEEE Computer, 31 (12):, December 1998, p. 55-61
25. Stichweh, R.: Systems Theory as an Alternative to Action Theory? The Rise of 'Communication' as a Theoretical Option. In: Acta Sociologica 43 (2000) p. 5-13
26. Wasserman, S., Faust, K.: Social Network Analysis. Methods and Applications. Cambridge University Press, Cambridge et al. (1994)
27. Wegner, P. Why interaction is more powerful than computing. Communications of the ACM, 40 (5):, May 1997. p. 80-91
28. Wooldridge, M.J., Jennings, N.R. and Kinny, D.. A methodology for agent-oriented analysis and design. In: Proceeding of the Third International Conference on Autonomous Agents.. ACM, Seattle (WA), May, 1-5 1999. p. 69-76